

## 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 Change History

Change history is described in the Rationale (Informative) volume of IEEE Std 1003.1-200x, and in the CHANGE HISTORY section of reference pages.

## 1.5 Terminology

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

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

### **can**

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

### **implementation-defined**

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

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

### **legacy**

Describes a feature or behavior that is being retained for compatibility with older applications, but which has limitations which make it inappropriate for developing portable

33 applications. New applications should use alternative means of obtaining equivalent  
34 functionality.

35 **may**

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

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

41 **shall**

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

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

46 **should**

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

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

53 **undefined**

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

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

60 **unspecified**

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

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

## 67 1.6 Definitions

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

## 69 1.7 Relationship to Other Formal Standards

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

72 ISO C (1999)

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

74 Parts of the ISO/IEC 9899: 1999 standard (hereinafter referred to as the ISO C standard) are  
75 referenced to describe requirements also mandated by this volume of IEEE Std 1003.1-200x.  
76 Some functions and headers included within this volume of IEEE Std 1003.1-200x have a version  
77 in the ISO C standard; in this case CX markings are added as appropriate to show where the  
78 ISO C standard has been extended (see Section 1.8.1). Any conflict between this volume of  
79 IEEE Std 1003.1-200x and the ISO C standard is unintentional.

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

## 82 1.8 Portability

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

87 Where optional, enhanced, or reduced functionality is specified, the text is shaded and a code in  
88 the margin identifies the nature of the option, extension, or warning (see Section 1.8.1). For  
89 maximum portability, an application should avoid such functionality.

### 90 1.8.1 Codes

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

93 ADV **Advisory Information**

94 The functionality described is optional. The functionality described is also an extension to the  
95 ISO C standard.

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

99 AIO **Asynchronous Input and Output**

100 The functionality described is optional. The functionality described is also an extension to the  
101 ISO C standard.

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

105 BAR **Barriers**

106 The functionality described is optional. The functionality described is also an extension to the

107 ISO C standard.

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

111 BE **Batch Environment Services and Utilities**  
112 The functionality described is optional.

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

116 CD **C-Language Development Utilities**  
117 The functionality described is optional.

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

121 CPT **Process CPU-Time Clocks**  
122 The functionality described is optional. The functionality described is also an extension to the  
123 ISO C standard.

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

127 CS **Clock Selection**  
128 The functionality described is optional. The functionality described is also an extension to the  
129 ISO C standard.

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

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

136 With each function or header from the ISO C standard, a statement to the effect that “any  
137 conflict is unintentional” is included. That is intended to refer to a direct conflict.  
138 IEEE Std 1003.1-200x acts in part as a profile of the ISO C standard, and it may choose to further  
139 constrain behaviors allowed to vary by the ISO C standard. Such limitations are not considered  
140 conflicts.

141 FD **FORTRAN Development Utilities**  
142 The functionality described is optional.

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

146 FR **FORTRAN Runtime Utilities**  
147 The functionality described is optional.

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

- 151 FSC **File Synchronization**  
152 The functionality described is optional. The functionality described is also an extension to the  
153 ISO C standard.
- 154 Where applicable, functions are marked with the FSC margin legend in the SYNOPSIS section.  
155 Where additional semantics apply to a function, the material is identified by use of the FSC  
156 margin legend.
- 157 IP6 **IPV6**  
158 The functionality described is optional. The functionality described is also an extension to the  
159 ISO C standard.
- 160 Where applicable, functions are marked with the IP6 margin legend in the SYNOPSIS section.  
161 Where additional semantics apply to a function, the material is identified by use of the IP6  
162 margin legend.
- 163 MC1 **Advisory Information and either Memory Mapped Files or Shared Memory Objects**  
164 The functionality described is optional. The functionality described is also an extension to the  
165 ISO C standard.
- 166 This is a shorthand notation for combinations of multiple option codes.
- 167 Where applicable, functions are marked with the MC1 margin legend in the SYNOPSIS section.  
168 Where additional semantics apply to a function, the material is identified by use of the MC1  
169 margin legend.
- 170 Refer to the Base Definitions volume of IEEE Std 1003.1-200x, Section 1.5.2, Margin Code  
171 Notation.
- 172 MC2 **Memory Mapped Files, Shared Memory Objects, or Memory Protection**  
173 The functionality described is optional. The functionality described is also an extension to the  
174 ISO C standard.
- 175 This is a shorthand notation for combinations of multiple option codes.
- 176 Where applicable, functions are marked with the MC2 margin legend in the SYNOPSIS section.  
177 Where additional semantics apply to a function, the material is identified by use of the MC2  
178 margin legend.
- 179 Refer to the Base Definitions volume of IEEE Std 1003.1-200x, Section 1.5.2, Margin Code  
180 Notation.
- 181 MF **Memory Mapped Files**  
182 The functionality described is optional. The functionality described is also an extension to the  
183 ISO C standard.
- 184 Where applicable, functions are marked with the MF margin legend in the SYNOPSIS section.  
185 Where additional semantics apply to a function, the material is identified by use of the MF  
186 margin legend.
- 187 ML **Process Memory Locking**  
188 The functionality described is optional. The functionality described is also an extension to the  
189 ISO C standard.
- 190 Where applicable, functions are marked with the ML margin legend in the SYNOPSIS section.  
191 Where additional semantics apply to a function, the material is identified by use of the ML  
192 margin legend.
- 193 MLR **Range Memory Locking**  
194 The functionality described is optional. The functionality described is also an extension to the

- 195 ISO C standard.
- 196 Where applicable, functions are marked with the MLR margin legend in the SYNOPSIS section.  
197 Where additional semantics apply to a function, the material is identified by use of the MLR  
198 margin legend.
- 199 MON **Monotonic Clock**  
200 The functionality described is optional. The functionality described is also an extension to the  
201 ISO C standard.
- 202 Where applicable, functions are marked with the MON margin legend in the SYNOPSIS section.  
203 Where additional semantics apply to a function, the material is identified by use of the MON  
204 margin legend.
- 205 MPR **Memory Protection**  
206 The functionality described is optional. The functionality described is also an extension to the  
207 ISO C standard.
- 208 Where applicable, functions are marked with the MPR margin legend in the SYNOPSIS section.  
209 Where additional semantics apply to a function, the material is identified by use of the MPR  
210 margin legend.
- 211 MSG **Message Passing**  
212 The functionality described is optional. The functionality described is also an extension to the  
213 ISO C standard.
- 214 Where applicable, functions are marked with the MSG margin legend in the SYNOPSIS section.  
215 Where additional semantics apply to a function, the material is identified by use of the MSG  
216 margin legend.
- 217 MX **IEC 60559 Floating-Point Option**  
218 The functionality described is optional. The functionality described is also an extension to the  
219 ISO C standard.
- 220 Where applicable, functions are marked with the MX margin legend in the SYNOPSIS section.  
221 Where additional semantics apply to a function, the material is identified by use of the MX  
222 margin legend.
- 223 OB **Obsolescent**  
224 The functionality described may be withdrawn in a future version of this volume of  
225 IEEE Std 1003.1-200x. Strictly Conforming POSIX Applications and Strictly Conforming XSI  
226 Applications shall not use obsolescent features.
- 227 OF **Output Format Incompletely Specified**  
228 The functionality described is an XSI extension. The format of the output produced by the utility  
229 is not fully specified. It is therefore not possible to post-process this output in a consistent  
230 fashion. Typical problems include unknown length of strings and unspecified field delimiters.
- 231 OH **Optional Header**  
232 In the SYNOPSIS section of some interfaces in the System Interfaces volume of  
233 IEEE Std 1003.1-200x an included header is marked as in the following example:
- 234 OH `#include <sys/types.h>`  
235 `#include <grp.h>`  
236 `struct group *getgrnam(const char *name);`
- 237 This indicates that the marked header is not required on XSI-conformant systems.

238	PIO	<b>Prioritized Input and Output</b>
239		The functionality described is optional. The functionality described is also an extension to the
240		ISO C standard.
241		Where applicable, functions are marked with the PIO margin legend in the SYNOPSIS section.
242		Where additional semantics apply to a function, the material is identified by use of the PIO
243		margin legend.
244	PS	<b>Process Scheduling</b>
245		The functionality described is optional. The functionality described is also an extension to the
246		ISO C standard.
247		Where applicable, functions are marked with the PS margin legend in the SYNOPSIS section.
248		Where additional semantics apply to a function, the material is identified by use of the PS
249		margin legend.
250	RS	<b>Raw Sockets</b>
251		The functionality described is optional. The functionality described is also an extension to the
252		ISO C standard.
253		Where applicable, functions are marked with the RS margin legend in the SYNOPSIS section.
254		Where additional semantics apply to a function, the material is identified by use of the RS
255		margin legend.
256	RTS	<b>Realtime Signals Extension</b>
257		The functionality described is optional. The functionality described is also an extension to the
258		ISO C standard.
259		Where applicable, functions are marked with the RTS margin legend in the SYNOPSIS section.
260		Where additional semantics apply to a function, the material is identified by use of the RTS
261		margin legend.
262	SD	<b>Software Development Utilities</b>
263		The functionality described is optional.
264		Where applicable, utilities are marked with the SD margin legend in the SYNOPSIS section.
265		Where additional semantics apply to a utility, the material is identified by use of the SD
266		margin legend.
267	SEM	<b>Semaphores</b>
268		The functionality described is optional. The functionality described is also an extension to the
269		ISO C standard.
270		Where applicable, functions are marked with the SEM margin legend in the SYNOPSIS section.
271		Where additional semantics apply to a function, the material is identified by use of the SEM
272		margin legend.
273	SHM	<b>Shared Memory Objects</b>
274		The functionality described is optional. The functionality described is also an extension to the
275		ISO C standard.
276		Where applicable, functions are marked with the SHM margin legend in the SYNOPSIS section.
277		Where additional semantics apply to a function, the material is identified by use of the SHM
278		margin legend.
279	SIO	<b>Synchronized Input and Output</b>
280		The functionality described is optional. The functionality described is also an extension to the
281		ISO C standard.

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

285 SPI **Spin Locks**

286 The functionality described is optional. The functionality described is also an extension to the  
287 ISO C standard.

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

291 SPN **Spawn**

292 The functionality described is optional. The functionality described is also an extension to the  
293 ISO C standard.

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

297 SS **Process Sporadic Server**

298 The functionality described is optional. The functionality described is also an extension to the  
299 ISO C standard.

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

303 TCT **Thread CPU-Time Clocks**

304 The functionality described is optional. The functionality described is also an extension to the  
305 ISO C standard.

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

309 TEF **Trace Event Filter**

310 The functionality described is optional. The functionality described is also an extension to the  
311 ISO C standard.

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

315 THR **Threads**

316 The functionality described is optional. The functionality described is also an extension to the  
317 ISO C standard.

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

321 TMO **Timeouts**

322 The functionality described is optional. The functionality described is also an extension to the  
323 ISO C standard.

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



327	TMR	<b>Timers</b>
328		The functionality described is optional. The functionality described is also an extension to the
329		ISO C standard.
330		Where applicable, functions are marked with the TMR margin legend in the SYNOPSIS section.
331		Where additional semantics apply to a function, the material is identified by use of the TMR
332		margin legend.
333	TPI	<b>Thread Priority Inheritance</b>
334		The functionality described is optional. The functionality described is also an extension to the
335		ISO C standard.
336		Where applicable, functions are marked with the TPI margin legend in the SYNOPSIS section.
337		Where additional semantics apply to a function, the material is identified by use of the TPI
338		margin legend.
339	TPP	<b>Thread Priority Protection</b>
340		The functionality described is optional. The functionality described is also an extension to the
341		ISO C standard.
342		Where applicable, functions are marked with the TPP margin legend in the SYNOPSIS section.
343		Where additional semantics apply to a function, the material is identified by use of the TPP
344		margin legend.
345	TPS	<b>Thread Execution Scheduling</b>
346		The functionality described is optional. The functionality described is also an extension to the
347		ISO C standard.
348		Where applicable, functions are marked with the TPS margin legend for the SYNOPSIS section.
349		Where additional semantics apply to a function, the material is identified by use of the TPS
350		margin legend.
351	TRC	<b>Trace</b>
352		The functionality described is optional. The functionality described is also an extension to the
353		ISO C standard.
354		Where applicable, functions are marked with the TRC margin legend in the SYNOPSIS section.
355		Where additional semantics apply to a function, the material is identified by use of the TRC
356		margin legend.
357	TRI	<b>Trace Inherit</b>
358		The functionality described is optional. The functionality described is also an extension to the
359		ISO C standard.
360		Where applicable, functions are marked with the TRI margin legend in the SYNOPSIS section.
361		Where additional semantics apply to a function, the material is identified by use of the TRI
362		margin legend.
363	TRL	<b>Trace Log</b>
364		The functionality described is optional. The functionality described is also an extension to the
365		ISO C standard.
366		Where applicable, functions are marked with the TRL margin legend in the SYNOPSIS section.
367		Where additional semantics apply to a function, the material is identified by use of the TRL
368		margin legend.
369	TSA	<b>Thread Stack Address Attribute</b>
370		The functionality described is optional. The functionality described is also an extension to the
371		ISO C standard.

372 Where applicable, functions are marked with the TSA margin legend for the SYNOPSIS section.  
373 Where additional semantics apply to a function, the material is identified by use of the TSA  
374 margin legend.

375 TSF **Thread-Safe Functions**

376 The functionality described is optional. The functionality described is also an extension to the  
377 ISO C standard.

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

381 TSH **Thread Process-Shared Synchronization**

382 The functionality described is optional. The functionality described is also an extension to the  
383 ISO C standard.

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

387 TSP **Thread Sporadic Server**

388 The functionality described is optional. The functionality described is also an extension to the  
389 ISO C standard.

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

393 TSS **Thread Stack Address Size**

394 The functionality described is optional. The functionality described is also an extension to the  
395 ISO C standard.

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

399 TYM **Typed Memory Objects**

400 The functionality described is optional. The functionality described is also an extension to the  
401 ISO C standard.

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

405 UP **User Portability Utilities**

406 The functionality described is optional.

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

410 XSI **Extension**

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

414 If an entire SYNOPSIS section is shaded and marked XSI, all the functionality described in that  
415 reference page is an extension. See the Base Definitions volume of IEEE Std 1003.1-200x, Section  
416 3.439, XSI.

417 XSR **XSI STREAMS**  
 418 The functionality described is optional. The functionality described is also an extension to the  
 419 ISO C standard.

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

## 423 1.9 Format of Entries

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

### 426 NAME

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

### 428 SYNOPSIS

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

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

### 433 DESCRIPTION

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

### 435 RETURN VALUE

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

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

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

### 444 ERRORS

445 This section gives the symbolic names of the error values returned by a function or  
 446 stored into a variable accessed through the symbol *errno* if an error occurs.

447 “No errors are defined” means that error values returned by a function or stored into a  
 448 variable accessed through the symbol *errno*, if any, depend on the implementation.

### 449 EXAMPLES

450 This section is non-normative.

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

### 454 APPLICATION USAGE

455 This section is non-normative.

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

- 459           **RATIONALE**
- 460           This section is non-normative.
- 461           This section contains historical information concerning the contents of this volume of
- 462           IEEE Std 1003.1-200x and why features were included or discarded by the standard
- 463           developers.
- 464           **FUTURE DIRECTIONS**
- 465           This section is non-normative.
- 466           This section provides comments which should be used as a guide to current thinking;
- 467           there is not necessarily a commitment to adopt these future directions.
- 468           **SEE ALSO**
- 469           This section is non-normative.
- 470           This section gives references to related information.
- 471           **CHANGE HISTORY**
- 472           This section is non-normative.
- 473           This section shows the derivation of the entry and any significant changes that have
- 474           been made to it. |

# General Information

475

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

## 478 2.1 Use and Implementation of Functions

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

- 481 1. If an argument to a function has an invalid value (such as a value outside the domain of  
 482 the function, or a pointer outside the address space of the program, or a null pointer), the  
 483 behavior is undefined.
- 484 2. Any function declared in a header may also be implemented as a macro defined in the  
 485 header, so a function should not be declared explicitly if its header is included. Any macro  
 486 definition of a function can be suppressed locally by enclosing the name of the function in  
 487 parentheses, because the name is then not followed by the left parenthesis that indicates  
 488 expansion of a macro function name. For the same syntactic reason, it is permitted to take  
 489 the address of a function even if it is also defined as a macro. The use of the C-language  
 490 `#undef` construct to remove any such macro definition shall also ensure that an actual  
 491 function is referred to.
- 492 3. Any invocation of a function that is implemented as a macro shall expand to code that  
 493 evaluates each of its arguments exactly once, fully protected by parentheses where  
 494 necessary, so it is generally safe to use arbitrary expressions as arguments. Likewise, those  
 495 function-like macros described in the following sections may be invoked in an expression  
 496 anywhere a function with a compatible return type could be called.
- 497 4. Provided that a function can be declared without reference to any type defined in a header,  
 498 it is also permissible to declare the function, either explicitly or implicitly, and use it  
 499 without including its associated header.
- 500 5. If a function that accepts a variable number of arguments is not declared (explicitly or by  
 501 including its associated header), the behavior is undefined.

## 502 2.2 The Compilation Environment

### 503 2.2.1 POSIX.1 Symbols

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

510 Symbols called “feature test macros” are used to control the visibility of symbols that might be  
 511 included in a header. Implementations, future versions of IEEE Std 1003.1-200x, and other |  
 512 standards may define additional feature test macros. |

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

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

#### 519 2.2.1.1 *The `_POSIX_C_SOURCE` Feature Test Macro*

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

522 When an application includes a header described by IEEE Std 1003.1-200x, and when this feature  
 523 test macro is defined to have the value `200xxxL`:

- 524 1. All symbols required by IEEE Std 1003.1-200x to appear when the header is included shall  
 525 be made visible.
- 526 2. Symbols that are explicitly permitted, but not required, by IEEE Std 1003.1-200x to appear  
 527 in that header (including those in reserved name spaces) may be made visible.
- 528 3. Additional symbols not required or explicitly permitted by IEEE Std 1003.1-200x to be in  
 529 that header shall not be made visible, except when enabled by another feature test macro.

530 Identifiers in IEEE Std 1003.1-200x may only be undefined using the **#undef** directive as  
 531 described in Section 2.1 (on page 463) or Section 2.2.2. These **#undef** directives shall follow all  
 532 **#include** directives of any header in IEEE Std 1003.1-200x.

533 **Note:** The POSIX.1-1990 standard specified a macro called `_POSIX_SOURCE`. This has been  
 534 superseded by `_POSIX_C_SOURCE`.

#### 535 2.2.1.2 *The `_XOPEN_SOURCE` Feature Test Macro*

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

540 Since this volume of IEEE Std 1003.1-200x is aligned with the ISO C standard, and since all  
 541 functionality enabled by `_POSIX_C_SOURCE` set equal to `200xxxL` is enabled by  
 542 `_XOPEN_SOURCE` set equal to `600`, there should be no need to define `_POSIX_C_SOURCE` if  
 543 `_XOPEN_SOURCE` is so defined. Therefore, if `_XOPEN_SOURCE` is set equal to `600` and  
 544 `_POSIX_C_SOURCE` is set equal to `200xxxL`, the behavior is the same as if only  
 545 `_XOPEN_SOURCE` is defined and set equal to `600`. However, should `_POSIX_C_SOURCE` be set  
 546 to a value greater than `200xxxL`, the behavior is unspecified.

### 547 2.2.2 The Name Space

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

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

- 560 1. Each identifier with external linkage described in the header section is reserved for use as  
561 an identifier with external linkage if the header is included.
- 562 2. Each macro described in the header section is reserved for any use if the header is  
563 included.
- 564 3. Each identifier with file scope described in the header section is reserved for use as an  
565 identifier with file scope in the same name space if the header is included.

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





619  
620  
621  
622  
623  
624  
625  
626  
627  
628  
629  
630  
631  
632  
633  
634  
635  
636  
637  
638

XSI  
  
  
XSI  
  
TMR  
TMR  
XSI  
XSI  
XSI

Header	Prefix	Suffix	Complete Name
<sys/uio.h>	iov_		UIO_MAXIOV
<sys/un.h>	sun_		
<sys/utsname.h>	uts_		
<sys/wait.h>	si_, W[A-Z], P_		
<termios.h>	c_		
<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	

639  
640  
641  
642

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

643  
644  
645  
646

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

647  
648  
649  
650  
651  
652  
653  
654  
655  
656  
657  
658  
659  
660  
661  
662  
663  
664  
665  
666  
667  
668  
669  
670  
671  
672  
673  
674  
675  
676  
677  
678  
679  
680  
681  
682  
683

Header	Prefix
<dlfcn.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_
	IPV6_, IN6_
<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/uioc.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_

684 **Note:** The notation [0-9] indicates any digit. The notation [A-Z] indicates any uppercase letter in the  
685 portable character set. The notation [0-9a-z\_] indicates any digit, any lowercase letter in the  
686 portable character set, or underscore.

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

688 XSI	CBAUD	EXTB	VDSUSP
689	DEFECHO	FLUSHO	VLNEXT
690	ECHOCTL	LOBLK	VREPRINT
691	ECHOKE	PENDIN	VSTATUS
692	ECHOPRT	SWTCH	VWERASE
693	EXTA	VDISCARD	

694 The following identifiers are reserved regardless of the inclusion of headers:

- 695 1. All identifiers that begin with an underscore and either an uppercase letter or another  
696 underscore are always reserved for any use by the implementation.
- 697 2. All identifiers that begin with an underscore are always reserved for use as identifiers with  
698 file scope in both the ordinary identifier and tag name spaces.
- 699 3. All identifiers in the table below are reserved for use as identifiers with external linkage.  
700 Some of these identifiers do not appear in this volume of IEEE Std 1003.1-200x, but are  
701 reserved for future use by the ISO C standard.

702	_Exit	cexp	fesetexceptflag	localtime	scalbn
703	abort	cexpf	fesetround	log	scalbnf
704	abs	cexpl	fetestexcept	log10	scalbnl
705	acos	cimag	feupdateenv	log10f	scanf
706	acosf	cimagf	fflush	log10l	setbuf
707	acosh	cimagl	fgetc	log1p	setjmp
708	acoshf	clearerr	fgetpos	log1pf	setlocale
709	acoshl	clock	fgets	log1pl	setvbuf
710	acosl	clog	fgetwc	log2	signal
711	acosl	clogf	fgetws	log2f	sin
712	asctime	clogl	floor	log2l	sinf
713	asin	conj	floorf	logb	sinh
714	asinf	conjf	floorl	logbf	sinhf
715	asinh	conjl	fma	logbl	sinhl
716	asinhf	copysign	fmaf	logf	sinl
717	asinhf	copysignf	fmal	logl	sprintf
718	asinl	copysignl	fmax	longjmp	sqrt
719	asinl	cos	fmaxf	lrint	sqrtf
720	atan	cosf	fmaxl	lrintf	sqrtl
721	atan2	cosh	fmin	lrintl	srand
722	atan2f	coshf	fminf	lround	sscanf
723	atan2l	coshl	fminl	lroundf	str[a-z]*
724	atanf	cosl	fmod	lroundl	strtof
725	atanf	cpow	fmodf	malloc	strtoimax
726	atanh	cpowf	fmodl	mblen	strtol
727	atanh	cpowl	fopen	mbrlen	strtoll
728	atanhf	cproj	fprintf	mbrtowc	strtoull
729	atanhl	cprojf	fputc	mbsinit	strtoumax
730	atanl	cprojl	fputs	mbsrtowcs	swprintf
731	atanl	creal	fputwc	mbstowcs	swscanf
732	atexit	crealf	fputws	mbtowc	system
733	atof	creall	fread	mem[a-z]*	tan
734	atoi	csin	free	mktime	tanf
735	atol	csinf	freopen	modf	tanh
736	atoll	csinh	frexp	modff	tanhf
737	bsearch	csinhf	frexpf	modfl	tanhf
738	cabs	csinhl	frexpl	nan	tanl
739	cabsf	csinl	fscanf	nanf	tgamma
740	cabsl	csqrt	fseek	nanl	tgammaf
741	cacos	csqrtf	fsetpos	nearbyint	tgammal

742	cacosf	csqrtl	ftell	nearbyintf	time
743	cacosh	ctan	fwide	nearbyintl	tmpfile
744	cacoshf	ctanf	fwprintf	nextafterf	tmpnam
745	cacoshl	ctanl	fwwrite	nextafterl	to[a-z]*
746	cacosl	ctime	fwscanf	nexttoward	trunc
747	calloc	difftime	getc	nexttowardf	truncf
748	carg	div	getchar	nexttowardl	truncl
749	cargf	erfcf	getenv	perror	ungetc
750	cargl	erfcl	gets	pow	ungetwc
751	casin	erff	getwc	powf	va_end
752	casinf	erfl	getwchar	powl	vfprintf
753	casinh	errno	gmtime	printf	vfscanf
754	casinhf	exit	hypotf	putc	vfwprintf
755	casinhl	exp	hypotl	putchar	vfwscanf
756	casinl	exp2	ilogb	puts	vprintf
757	catan	exp2f	ilogbf	putwc	vscanf
758	catanf	exp2l	ilogbl	putwchar	vsprintf
759	catanh	expf	imaxabs	qsort	vsscanf
760	catanh	expl	imaxdiv	raise	vswprintf
761	catanhf	expm1	is[a-z]*	rand	vswscanf
762	catanhf	expm1f	isblank	realloc	vwprintf
763	catanhl	expm1l	iswblank	remainderf	vwscanf
764	catanhl	fabs	labs	remainderl	wcrtomb
765	catanl	fabsf	ldexp	remove	wcs[a-z]*
766	cbrt	fabsl	ldexpf	remquo	wcstof
767	cbrtf	fclose	ldexpl	remquof	wcstoimax
768	cbrtl	fdim	ldiv	remquol	wcstold
769	ccos	fdimf	ldiv	rename	wcstoll
770	ccosf	fdiml	lgammaf	rewind	wcstoull
771	ccosh	feclearexcept	lgammal	rint	wcstoumax
772	ccoshf	fegetenv	llabs	rintf	wctob
773	ccoshl	fegetexceptflag	llrint	rintl	wctomb
774	ccosl	fegetround	llrintf	round	wctrans
775	ceil	feholdexcept	llrintl	roundf	wctype
776	ceilf	feof	llround	roundl	wcwidth
777	ceilf	feraiseexcept	llroundf	scalbln	wmem[a-z]*
778	ceil	ferror	llroundl	scalblnf	wprintf
779	ceil	fesetenv	localeconv	scalblnl	wscanf

780 **Note:** The notation [a-z] indicates any lowercase letter in the portable character set. The  
781 notation ' \* ' indicates any combination of digits, letters in the portable character set, or  
782 underscore.

783 4. All functions and external identifiers defined in the Base Definitions volume of  
784 IEEE Std 1003.1-200x, Chapter 13, Headers are reserved for use as identifiers with external  
785 linkage.

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

788 No other identifiers are reserved.

789 Applications shall not declare or define identifiers with the same name as an identifier reserved  
790 in the same context. Since macro names are replaced whenever found, independent of scope and  
791 name space, macro names matching any of the reserved identifier names shall not be defined by

792 an application if any associated header is included.

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

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

## 802 **2.3 Error Numbers**

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

805 Some functions provide the error number in a variable accessed through the symbol *errno*. The  
806 symbol *errno*, defined by including the `<errno.h>` header, expands to a modifiable lvalue of type  
807 `int`. It is unspecified whether *errno* is a macro or an identifier declared with external linkage. If a  
808 macro definition is suppressed in order to access an actual object, or a program defines an  
809 identifier with the name *errno*, the behavior is undefined.

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

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

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

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

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

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

831 The following symbolic names identify the possible error numbers, in the context of the  
832 functions specifically defined in this volume of IEEE Std 1003.1-200x; these general descriptions  
833 are more precisely defined in the ERRORS sections of the functions that return them. Only these  
834 symbolic names should be used in programs, since the actual value of the error number is  
835 unspecified. All values listed in this section shall be unique integer constant expressions with

836 type **int** suitable for use in **#if** preprocessing directives, except as noted below. The values for all  
 837 these names shall be found in the **<errno.h>** header defined in the Base Definitions volume of  
 838 IEEE Std 1003.1-200x. The actual values are unspecified by this volume of IEEE Std 1003.1-200x.

839 [E2BIG]

840 Argument list too long. The sum of the number of bytes used by the new process image's  
 841 argument list and environment list is greater than the system-imposed limit of {ARG\_MAX}  
 842 bytes.

843 or:

844 Lack of space in an output buffer.

845 or:

846 Argument is greater than the system-imposed maximum.

847 [EACCES]

848 Permission denied. An attempt was made to access a file in a way forbidden by its file  
 849 access permissions.

850 [EADDRINUSE]

851 Address in use. The specified address is in use.

852 [EADDRNOTAVAIL]

853 Address not available. The specified address is not available from the local system.

854 [EAFNOSUPPORT]

855 Address family not supported. The implementation does not support the specified address  
 856 family, or the specified address is not a valid address for the address family of the specified  
 857 socket.

858 [EAGAIN]

859 Resource temporarily unavailable. This is a temporary condition and later calls to the same  
 860 routine may complete normally.

861 [EALREADY]

862 Connection already in progress. A connection request is already in progress for the specified  
 863 socket.

864 [EBADF]

865 Bad file descriptor. A file descriptor argument is out of range, refers to no open file, or a  
 866 read (write) request is made to a file that is only open for writing (reading).

867 [EBADMSG]

868 XSR Bad message. During a *read()*, *getmsg()*, *getpmsg()*, or *ioctl()* I\_RECVFD request to a |  
 869 STREAMS device, a message arrived at the head of the STREAM that is inappropriate for |  
 870 the function receiving the message.

871 *read()* Message waiting to be read on a STREAM is not a data message. |

872 *getmsg()* or *getpmsg()* |

873 A file descriptor was received instead of a control message. |

874 *ioctl()* Control or data information was received instead of a file descriptor when |  
 875 I\_RECVFD was specified.

876 or:

877 Bad Message. The implementation has detected a corrupted message.

878	[EBUSY]
879	Resource busy. An attempt was made to make use of a system resource that is not currently
880	available, as it is being used by another process in a manner that would have conflicted with
881	the request being made by this process.
882	[ECANCELED]
883	Operation canceled. The associated asynchronous operation was canceled before
884	completion.
885	[ECHILD]
886	No child process. A <i>wait()</i> or <i>waitpid()</i> function was executed by a process that had no
887	existing or unwaited-for child process.
888	[ECONNABORTED]
889	Connection aborted. The connection has been aborted.
890	[ECONNREFUSED]
891	Connection refused. An attempt to connect to a socket was refused because there was no
892	process listening or because the queue of connection requests was full and the underlying
893	protocol does not support retransmissions.
894	[ECONNRESET]
895	Connection reset. The connection was forcibly closed by the peer.
896	[EDEADLK]
897	Resource deadlock would occur. An attempt was made to lock a system resource that
898	would have resulted in a deadlock situation.
899	[EDESTADDRREQ]
900	Destination address required. No bind address was established.
901	[EDOM]
902	Domain error. An input argument is outside the defined domain of the mathematical
903	function (defined in the ISO C standard).
904	[EDQUOT]
905	Reserved.
906	[EEXIST]
907	File exists. An existing file was mentioned in an inappropriate context; for example, as a
908	new link name in the <i>link()</i> function.
909	[EFAULT]
910	Bad address. The system detected an invalid address in attempting to use an argument of a
911	call. The reliable detection of this error cannot be guaranteed, and when not detected may
912	result in the generation of a signal, indicating an address violation, which is sent to the
913	process.
914	[EFBIG]
915	File too large. The size of a file would exceed the maximum file size of an implementation or
916	offset maximum established in the corresponding file description.
917	[EHOSTUNREACH]
918	Host is unreachable. The destination host cannot be reached (probably because the host is
919	down or a remote router cannot reach it).
920	[EIDRM]
921	Identifier removed. Returned during XSI interprocess communication if an identifier has
922	been removed from the system.

923	[EILSEQ]
924	Illegal byte sequence. A wide-character code has been detected that does not correspond to
925	a valid character, or a byte sequence does not form a valid wide-character code (defined in
926	the ISO C standard).
927	[EINPROGRESS]
928	Operation in progress. This code is used to indicate that an asynchronous operation has not
929	yet completed.
930	or:
931	O_NONBLOCK is set for the socket file descriptor and the connection cannot be
932	immediately established.
933	[EINTR]
934	Interrupted function call. An asynchronous signal was caught by the process during the
935	execution of an interruptible function. If the signal handler performs a normal return, the
936	interrupted function call may return this condition (see the Base Definitions volume of
937	IEEE Std 1003.1-200x, <signal.h>).
938	[EINVAL]
939	Invalid argument. Some invalid argument was supplied; for example, specifying an
940	undefined signal in a <i>signal()</i> function or a <i>kill()</i> function.
941	[EIO]
942	Input/output error. Some physical input or output error has occurred. This error may be
943	reported on a subsequent operation on the same file descriptor. Any other error-causing
944	operation on the same file descriptor may cause the [EIO] error indication to be lost.
945	[EISCONN]
946	Socket is connected. The specified socket is already connected.
947	[EISDIR]
948	Is a directory. An attempt was made to open a directory with write mode specified.
949	[ELOOP]
950	Symbolic link loop. A loop exists in symbolic links encountered during pathname
951	resolution. This error may also be returned if more than {SYMLOOP_MAX} symbolic links
952	are encountered during pathname resolution.
953	[EMFILE]
954	Too many open files. An attempt was made to open more than the maximum number of
955	{OPEN_MAX} file descriptors allowed in this process.
956	[EMLINK]
957	Too many links. An attempt was made to have the link count of a single file exceed
958	{LINK_MAX}.
959	[EMSGSIZE]
960	Message too large. A message sent on a transport provider was larger than an internal
961	message buffer or some other network limit.
962	or:
963	Inappropriate message buffer length.
964	[EMULTIHOP]
965	Reserved.



966	[ENAMETOOLONG]	
967	Filename too long. The length of a pathname exceeds {PATH_MAX}, or a pathname	
968	component is longer than {NAME_MAX}. This error may also occur when pathname	
969	substitution, as a result of encountering a symbolic link during pathname resolution, results	
970	in a pathname string the size of which exceeds {PATH_MAX}.	
971	[ENETDOWN]	
972	Network is down. The local network interface used to reach the destination is down.	
973	[ENETRESET]	
974	The connection was aborted by the network.	
975	[ENETUNREACH]	
976	Network unreachable. No route to the network is present.	
977	[ENFILE]	
978	Too many files open in system. Too many files are currently open in the system. The system	
979	has reached its predefined limit for simultaneously open files and temporarily cannot accept	
980	requests to open another one.	
981	[ENOBUFS]	
982	No buffer space available. Insufficient buffer resources were available in the system to	
983	perform the socket operation.	
984	XSR [ENODATA]	
985	No message available. No message is available on the STREAM head read queue.	
986	[ENODEV]	
987	No such device. An attempt was made to apply an inappropriate function to a device; for	
988	example, trying to read a write-only device such as a printer.	
989	[ENOENT]	
990	No such file or directory. A component of a specified pathname does not exist, or the	
991	pathname is an empty string.	
992	[ENOEXEC]	
993	Executable file format error. A request is made to execute a file that, although it has the	
994	appropriate permissions, is not in the format required by the implementation for executable	
995	files.	
996	[ENOLCK]	
997	No locks available. A system-imposed limit on the number of simultaneous file and record	
998	locks has been reached and no more are currently available.	
999	[ENOLINK]	
1000	Reserved.	
1001	[ENOMEM]	
1002	Not enough space. The new process image requires more memory than is allowed by the	
1003	hardware or system-imposed memory management constraints.	
1004	[ENOMSG]	
1005	No message of the desired type. The message queue does not contain a message of the	
1006	required type during XSI interprocess communication.	
1007	[ENOPROTOPT]	
1008	Protocol not available. The protocol option specified to <i>setsockopt()</i> is not supported by the	
1009	implementation.	

1010	[ENOSPC]	
1011		No space left on a device. During the <i>write()</i> function on a regular file or when extending a
1012		directory, there is no free space left on the device.
1013	XSR [ENOSR]	
1014		No STREAM resources. Insufficient STREAMS memory resources are available to perform a
1015		STREAMS-related function. This is a temporary condition; it may be recovered from if other
1016		processes release resources.
1017	XSR [ENOSTR]	
1018		Not a STREAM. A STREAM function was attempted on a file descriptor that was not
1019		associated with a STREAMS device.
1020	[ENOSYS]	
1021		Function not implemented. An attempt was made to use a function that is not available in
1022		this implementation.
1023	[ENOTCONN]	
1024		Socket not connected. The socket is not connected.
1025	[ENOTDIR]	
1026		Not a directory. A component of the specified pathname exists, but it is not a directory,
1027		when a directory was expected.
1028	[ENOTEMPTY]	
1029		Directory not empty. A directory other than an empty directory was supplied when an
1030		empty directory was expected.
1031	[ENOTSOCK]	
1032		Not a socket. The file descriptor does not refer to a socket.
1033	[ENOTSUP]	
1034		Not supported. The implementation does not support this feature of the Realtime Option
1035		Group.
1036	[ENOTTY]	
1037		Inappropriate I/O control operation. A control function has been attempted for a file or
1038		special file for which the operation is inappropriate.
1039	[ENXIO]	
1040		No such device or address. Input or output on a special file refers to a device that does not
1041		exist, or makes a request beyond the capabilities of the device. It may also occur when, for
1042		example, a tape drive is not on-line.
1043	[EOPNOTSUPP]	
1044		Operation not supported on socket. The type of socket (address family or protocol) does not
1045		support the requested operation.
1046	[EOVERFLOW]	
1047		Value too large to be stored in data type. An operation was attempted which would
1048		generate a value that is outside the range of values that can be represented in the relevant
1049		data type or that are allowed for a given data item.
1050	[EPERM]	
1051		Operation not permitted. An attempt was made to perform an operation limited to
1052		processes with appropriate privileges or to the owner of a file or other resource.
1053	[EPIPE]	
1054		Broken pipe. A write was attempted on a socket, pipe, or FIFO for which there is no process

1055		to read the data.
1056		[EPROTO]
1057		Protocol error. Some protocol error occurred. This error is device-specific, but is generally
1058		not related to a hardware failure.
1059		[EPROTONOSUPPORT]
1060		Protocol not supported. The protocol is not supported by the address family, or the protocol
1061		is not supported by the implementation.
1062		[EPROTOTYPE]
1063		Protocol wrong type for socket. The socket type is not supported by the protocol.
1064		[ERANGE]
1065		Result too large or too small. The result of the function is too large (overflow) or too small
1066		(underflow) to be represented in the available space (defined in the ISO C standard).
1067		[EROFS]
1068		Read-only file system. An attempt was made to modify a file or directory on a file system
1069		that is read-only.
1070		[ESPIPE]
1071		Invalid seek. An attempt was made to access the file offset associated with a pipe or FIFO.
1072		[ESRCH]
1073		No such process. No process can be found corresponding to that specified by the given
1074		process ID.
1075		[ESTALE]
1076		Reserved.
1077	XSR	[ETIME]
1078		STREAM <i>ioctl()</i> timeout. The timer set for a STREAMS <i>ioctl()</i> call has expired. The cause of
1079		this error is device-specific and could indicate either a hardware or software failure, or a
1080		timeout value that is too short for the specific operation. The status of the <i>ioctl()</i> operation
1081		is unspecified.
1082		[ETIMEDOUT]
1083		Connection timed out. The connection to a remote machine has timed out. If the connection
1084		timed out during execution of the function that reported this error (as opposed to timing
1085		out prior to the function being called), it is unspecified whether the function has completed
1086		some or all of the documented behavior associated with a successful completion of the
1087		function.
1088		or:
1089		Operation timed out. The time limit associated with the operation was exceeded before the
1090		operation completed.
1091		[ETXTBSY]
1092		Text file busy. An attempt was made to execute a pure-procedure program that is currently
1093		open for writing, or an attempt has been made to open for writing a pure-procedure
1094		program that is being executed.
1095		[EWOULDBLOCK]
1096		Operation would block. An operation on a socket marked as non-blocking has encountered
1097		a situation such as no data available that otherwise would have caused the function to
1098		suspend execution.

1099 A conforming implementation may assign the same values for [EWOULDBLOCK] and  
1100 [EAGAIN].  
1101 [EXDEV]  
1102 Improper link. A link to a file on another file system was attempted.

### 1103 2.3.1 Additional Error Numbers

1104 Additional implementation-defined error numbers may be defined in `<errno.h>`.

## 1105 2.4 Signal Concepts

### 1106 2.4.1 Signal Generation and Delivery

1107 A signal is said to be *generated* for (or sent to) a process or thread when the event that causes the  
1108 signal first occurs. Examples of such events include detection of hardware faults, timer  
1109 RTS expiration, signals generated via the `sigevent` structure and terminal activity, as well as  
1110 RTS invocations of the `kill()` and `sigqueue()` functions. In some circumstances, the same event  
1111 generates signals for multiple processes.

1112 At the time of generation, a determination shall be made whether the signal has been generated  
1113 for the process or for a specific thread within the process. Signals which are generated by some  
1114 action attributable to a particular thread, such as a hardware fault, shall be generated for the  
1115 thread that caused the signal to be generated. Signals that are generated in association with a  
1116 process ID or process group ID or an asynchronous event, such as terminal activity, shall be  
1117 generated for the process.

1118 Each process has an action to be taken in response to each signal defined by the system (see  
1119 Section 2.4.3 (on page 480)). A signal is said to be *delivered* to a process when the appropriate  
1120 action for the process and signal is taken. A signal is said to be *accepted* by a process when the  
1121 signal is selected and returned by one of the `sigwait()` functions.

1122 During the time between the generation of a signal and its delivery or acceptance, the signal is  
1123 said to be *pending*. Ordinarily, this interval cannot be detected by an application. However, a  
1124 signal can be *blocked* from delivery to a thread. If the action associated with a blocked signal is  
1125 anything other than to ignore the signal, and if that signal is generated for the thread, the signal  
1126 shall remain pending until it is unblocked, it is accepted when it is selected and returned by a  
1127 call to the `sigwait()` function, or the action associated with it is set to ignore the signal. Signals  
1128 generated for the process shall be delivered to exactly one of those threads within the process  
1129 which is in a call to a `sigwait()` function selecting that signal or has not blocked delivery of the  
1130 signal. If there are no threads in a call to a `sigwait()` function selecting that signal, and if all  
1131 threads within the process block delivery of the signal, the signal shall remain pending on the  
1132 process until a thread calls a `sigwait()` function selecting that signal, a thread unblocks delivery  
1133 of the signal, or the action associated with the signal is set to ignore the signal. If the action  
1134 associated with a blocked signal is to ignore the signal and if that signal is generated for the  
1135 process, it is unspecified whether the signal is discarded immediately upon generation or  
1136 remains pending.

1137 Each thread has a *signal mask* that defines the set of signals currently blocked from delivery to it. |  
1138 The signal mask for a thread shall be initialized from that of its parent or creating thread, or from |  
1139 the corresponding thread in the parent process if the thread was created as the result of a call to |  
1140 `fork()`. The `sigaction()`, `sigprocmask()`, and `sigsuspend()` functions control the manipulation of the |  
1141 signal mask.

1142 The determination of which action is taken in response to a signal is made at the time the signal  
 1143 is delivered, allowing for any changes since the time of generation. This determination is  
 1144 independent of the means by which the signal was originally generated. If a subsequent  
 1145 occurrence of a pending signal is generated, it is implementation-defined as to whether the  
 1146 RTS signal is delivered or accepted more than once in circumstances other than those in which  
 1147 queuing is required under the Realtime Signals Extension option. The order in which multiple,  
 1148 simultaneously pending signals outside the range SIGRTMIN to SIGRTMAX are delivered to or  
 1149 accepted by a process is unspecified.

1150 When any stop signal (SIGSTOP, SIGTSTP, SIGTTIN, SIGTTOU) is generated for a process, any  
 1151 pending SIGCONT signals for that process shall be discarded. Conversely, when SIGCONT is  
 1152 generated for a process, all pending stop signals for that process shall be discarded. When  
 1153 SIGCONT is generated for a process that is stopped, the process shall be continued, even if the  
 1154 SIGCONT signal is blocked or ignored. If SIGCONT is blocked and not ignored, it shall remain  
 1155 pending until it is either unblocked or a stop signal is generated for the process.

1156 An implementation shall document any condition not specified by this volume of  
 1157 IEEE Std 1003.1-200x under which the implementation generates signals.

1158 **2.4.2 Realtime Signal Generation and Delivery**

1159 RTS This section describes extensions to support realtime signal generation and delivery. This  
 1160 functionality is dependent on support of the Realtime Signals Extension option (and the rest of  
 1161 this section is not further shaded for this option).

1162 Some signal-generating functions, such as high-resolution timer expiration, asynchronous I/O  
 1163 completion, interprocess message arrival, and the *sigqueue()* function, support the specification  
 1164 of an application-defined value, either explicitly as a parameter to the function or in a **sigevent**  
 1165 structure parameter. The **sigevent** structure is defined in **<signal.h>** and contains at least the  
 1166 following members:

1167

1168

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

1174 The *sigev\_notify* member specifies the notification mechanism to use when an asynchronous  
 1175 event occurs. This volume of IEEE Std 1003.1-200x defines the following values for the  
 1176 *sigev\_notify* member:

1177 SIGEV\_NONE No asynchronous notification shall be delivered when the event of  
 1178 interest occurs.

1179 SIGEV\_SIGNAL The signal specified in *sigev\_signo* shall be generated for the process when  
 1180 the event of interest occurs. If the implementation supports the Realtime  
 1181 Signals Extension option and if the SA\_SIGINFO flag is set for that signal  
 1182 number, then the signal shall be queued to the process and the value  
 1183 specified in *sigev\_value* shall be the *si\_value* component of the generated  
 1184 signal. If SA\_SIGINFO is not set for that signal number, it is unspecified  
 1185 whether the signal is queued and what value, if any, is sent.

1186 SIGEV\_THREAD A notification function shall be called to perform notification.

1187 An implementation may define additional notification mechanisms.

1188 The *sigev\_signo* member specifies the signal to be generated. The *sigev\_value* member is the  
 1189 application-defined value to be passed to the signal-catching function at the time of the signal  
 1190 delivery or to be returned at signal acceptance as the *si\_value* member of the **siginfo\_t** structure.

1191 The **signal** union is defined in `<signal.h>` and contains at least the following members:

1192

1193

1194

1195

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

1196 The *sival\_int* member shall be used when the application-defined value is of type **int**; the  
 1197 *sival\_ptr* member shall be used when the application-defined value is a pointer.

1198 When a signal is generated by the *sigqueue()* function or any signal-generating function that  
 1199 supports the specification of an application-defined value, the signal shall be marked pending  
 1200 and, if the SA\_SIGINFO flag is set for that signal, the signal shall be queued to the process along  
 1201 with the application-specified signal value. Multiple occurrences of signals so generated are  
 1202 queued in FIFO order. It is unspecified whether signals so generated are queued when the  
 1203 SA\_SIGINFO flag is not set for that signal.

1204 Signals generated by the *kill()* function or other events that cause signals to occur, such as  
 1205 detection of hardware faults, *alarm()* timer expiration, or terminal activity, and for which the  
 1206 implementation does not support queuing, shall have no effect on signals already queued for the  
 1207 same signal number.

1208 When multiple unblocked signals, all in the range SIGRTMIN to SIGRTMAX, are pending, the  
 1209 behavior shall be as if the implementation delivers the pending unblocked signal with the lowest  
 1210 signal number within that range. No other ordering of signal delivery is specified.

1211 If, when a pending signal is delivered, there are additional signals queued to that signal number,  
 1212 the signal shall remain pending. Otherwise, the pending indication shall be reset.

1213 Multi-threaded programs can use an alternate event notification mechanism. When a  
 1214 notification is processed, and the *sigev\_notify* member of the **sigevent** structure has the value  
 1215 SIGEV\_THREAD, the function *sigev\_notify\_function* is called with parameter *sigev\_value*.

1216 The function shall be executed in an environment as if it were the *start\_routine* for a newly  
 1217 created thread with thread attributes specified by *sigev\_notify\_attributes*. If *sigev\_notify\_attributes*  
 1218 is NULL, the behavior shall be as if the thread were created with the *detachstate* attribute set to  
 1219 PTHREAD\_CREATE\_DETACHED. Supplying an attributes structure with a *detachstate* attribute  
 1220 of PTHREAD\_CREATE\_JOINABLE results in undefined behavior. The signal mask of this  
 1221 thread is implementation-defined.

### 1222 2.4.3 Signal Actions

1223 There are three types of action that can be associated with a signal: SIG\_DFL, SIG\_IGN, or a  
 1224 pointer to a function. Initially, all signals shall be set to SIG\_DFL or SIG\_IGN prior to entry of  
 1225 the *main()* routine (see the *exec* functions). The actions prescribed by these values are as follows:

1226 SIG\_DFL Signal-specific default action.

1227 The default actions for the signals defined in this volume of IEEE Std 1003.1-200x  
 1228 RTS are specified under `<signal.h>`. If the Realtime Signals Extension option is  
 1229 supported, the default actions for the realtime signals in the range SIGRTMIN to  
 1230 SIGRTMAX shall be to terminate the process abnormally.

1231		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
1232		generated for its parent process, unless the parent process has set the
1233		SA_NOCLDSTOP flag. While a process is stopped, any additional signals that are
1234		sent to the process shall not be delivered until the process is continued, except
1235		SIGKILL which always terminates the receiving process. A process that is a
1236		member of an orphaned process group shall not be allowed to stop in response to
1237		the SIGTSTP, SIGTTIN, or SIGTTOU signals. In cases where delivery of one of
1238		these signals would stop such a process, the signal shall be discarded.
1239		
1240		Setting a signal action to SIG_DFL for a signal that is pending, and whose default
1241		action is to ignore the signal (for example, SIGCHLD), shall cause the pending
1242		signal to be discarded, whether or not it is blocked.
1243		The default action for SIGCONT is to resume execution at the point where the
1244	RTS	process was stopped, after first handling any pending unblocked signals. If the
1245		Realtime Signals Extension option is supported, any queued values pending shall
1246		be discarded and the resources used to queue them shall be released and returned
1247		to the system for other use. When a stopped process is continued, a SIGCHLD
1248		signal may be generated for its parent process, unless the parent process has set
1249		the SA_NOCLDSTOP flag.
1250	SIG_IGN	Ignore signal.
1251		Delivery of the signal shall have no effect on the process. The behavior of a process
1252	RTS	is undefined after it ignores a SIGFPE, SIGILL, SIGSEGV, or SIGBUS signal that
1253	RTS	was not generated by <i>kill()</i> , <i>sigqueue()</i> , or <i>raise()</i> .
1254		The system shall not allow the action for the signals SIGKILL or SIGSTOP to be set
1255		to SIG_IGN.
1256		Setting a signal action to SIG_IGN for a signal that is pending shall cause the
1257		pending signal to be discarded, whether or not it is blocked.
1258		If a process sets the action for the SIGCHLD signal to SIG_IGN, the behavior is
1259	XSI	unspecified, except as specified below.
1260		If the action for the SIGCHLD signal is set to SIG_IGN, child processes of the
1261		calling processes shall not be transformed into zombie processes when they
1262		terminate. If the calling process subsequently waits for its children, and the process
1263		has no unwaited-for children that were transformed into zombie processes, it shall
1264		block until all of its children terminate, and <i>wait()</i> , <i>waitid()</i> , and <i>waitpid()</i> shall fail
1265		and set <i>errno</i> to [ECHILD].
1266	RTS	If the Realtime Signals Extension option is supported, any queued values pending
1267		shall be discarded and the resources used to queue them shall be released and
1268		made available to queue other signals.
1269		<i>pointer to a function</i>
1270		Catch signal.
1271		On delivery of the signal, the receiving process is to execute the signal-catching
1272		function at the specified address. After returning from the signal-catching function,
1273		the receiving process shall resume execution at the point at which it was
1274		interrupted.
1275		If the SA_SIGINFO flag for the signal is cleared, the signal-catching function shall
1276		be entered as a C-language function call as follows:

1277 `void func(int signo);`

1278 XSI|RTS If the SA\_SIGINFO flag for the signal is set, the signal-catching function shall be  
1279 entered as a C-language function call as follows:

1280 `void func(int signo, siginfo_t *info, void *context);`

1281 where *func* is the specified signal-catching function, *signo* is the signal number of  
1282 the signal being delivered, and *info* is a pointer to a **siginfo\_t** structure defined in  
1283 **<signal.h>** containing at least the following members:

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.

1289 The *si\_signo* member shall contain the signal number. This shall be the same as the  
1290 *signo* parameter. The *si\_code* member shall contain a code identifying the cause of  
1291 the signal. The following values are defined for *si\_code*:

1292 SI\_USER The signal was sent by the *kill()* function. The implementation  
1293 may set *si\_code* to SI\_USER if the signal was sent by the *raise()* or  
1294 *abort()* functions or any similar functions provided as  
1295 implementation extensions.

1296 RTS SI\_QUEUE The signal was sent by the *sigqueue()* function.

1297 RTS SI\_TIMER The signal was generated by the expiration of a timer set by  
1298 *timer\_settime()*.

1299 RTS SI\_ASYNCIO The signal was generated by the completion of an asynchronous  
1300 I/O request.

1301 RTS SI\_MESGQ The signal was generated by the arrival of a message on an  
1302 empty message queue.

1303 If the signal was not generated by one of the functions or events listed above, the  
1304 *si\_code* shall be set to an implementation-defined value that is not equal to any of  
1305 the values defined above.

1306 RTS If the Realtime Signals Extension is supported, and *si\_code* is one of SI\_QUEUE,  
1307 SI\_TIMER, SI\_ASYNCIO, or SI\_MESGQ, then *si\_value* shall contain the  
1308 application-specified signal value. Otherwise, the contents of *si\_value* are  
1309 undefined.

1310 The behavior of a process is undefined after it returns normally from a signal-  
1311 XSI catching function for a SIGBUS, SIGFPE, SIGILL, or SIGSEGV signal that was not  
1312 RTS generated by *kill()*, *sigqueue()*, or *raise()*.

1313 The system shall not allow a process to catch the signals SIGKILL and SIGSTOP.

1314 If a process establishes a signal-catching function for the SIGCHLD signal while it  
1315 has a terminated child process for which it has not waited, it is unspecified  
1316 whether a SIGCHLD signal is generated to indicate that child process.

1317 When signal-catching functions are invoked asynchronously with process  
1318 execution, the behavior of some of the functions defined by this volume of  
1319 IEEE Std 1003.1-200x is unspecified if they are called from a signal-catching  
1320 function.



1321 The following table defines a set of functions that shall be either reentrant or non-  
 1322 interruptible by signals and shall be async-signal-safe. Therefore applications may  
 1323 invoke them, without restriction, from signal-catching functions:

### 1324 **Notes to Reviewers**

1325 *This section with side shading will not appear in the final copy. - Ed.*

1326 The contents of the following tables need to be reviewed.

1327	<code>_Exit()</code>	<code>fdatasync()</code>	<code>posix_trace_event()</code>	<code>sigsuspend()</code>
1328	<code>_exit()</code>	<code>fork()</code>	<code>raise()</code>	<code>stat()</code>
1329	<code>access()</code>	<code>fpathconf()</code>	<code>read()</code>	<code>symlink()</code>
1330	<code>aio_error()</code>	<code>fstat()</code>	<code>readlink()</code>	<code>sysconf()</code>
1331	<code>aio_return()</code>	<code>fsync()</code>	<code>rename()</code>	<code>tcdrain()</code>
1332	<code>aio_suspend()</code>	<code>ftruncate()</code>	<code>rmdir()</code>	<code>tcflow()</code>
1333	<code>alarm()</code>	<code>getegid()</code>	<code>sem_post()</code>	<code>tcflush()</code>
1334	<code>cfgetispeed()</code>	<code>geteuid()</code>	<code>setgid()</code>	<code>tcgetattr()</code>
1335	<code>cfgetospeed()</code>	<code>getgid()</code>	<code>setpgid()</code>	<code>tcgetpgrp()</code>
1336	<code>cfsetispeed()</code>	<code>getgroups()</code>	<code>setsid()</code>	<code>tcsendbreak()</code>
1337	<code>cfsetospeed()</code>	<code>getpgrp()</code>	<code>setuid()</code>	<code>tcsetattr()</code>
1338	<code>chdir()</code>	<code>getpid()</code>	<code>sigaction()</code>	<code>tcsetpgrp()</code>
1339	<code>chmod()</code>	<code>getppid()</code>	<code>sigaddset()</code>	<code>time()</code>
1340	<code>chown()</code>	<code>getuid()</code>	<code>sigdelset()</code>	<code>timer_getoverrun()</code>
1341	<code>clock_gettime()</code>	<code>kill()</code>	<code>sigemptyset()</code>	<code>timer_gettime()</code>
1342	<code>close()</code>	<code>link()</code>	<code>sigfillset()</code>	<code>timer_settime()</code>
1343	<code>creat()</code>	<code>lseek()</code>	<code>sigismember()</code>	<code>times()</code>
1344	<code>dup()</code>	<code>lstat()</code>	<code>sleep()</code>	<code>umask()</code>
1345	<code>dup2()</code>	<code>mkdir()</code>	<code>signal()</code>	<code>uname()</code>
1346	<code>execle()</code>	<code>mkfifo()</code>	<code>sigpause()</code>	<code>unlink()</code>
1347	<code>execve()</code>	<code>open()</code>	<code>sigpending()</code>	<code>utime()</code>
1348	<code>fchmod()</code>	<code>pathconf()</code>	<code>sigprocmask()</code>	<code>wait()</code>
1349	<code>fchown()</code>	<code>pause()</code>	<code>sigqueue()</code>	<code>waitpid()</code>
1350	<code>fcntl()</code>	<code>pipe()</code>	<code>sigset()</code>	<code>write()</code>

1351 All functions not in the above table are considered to be unsafe with respect to  
 1352 signals. In the presence of signals, all functions defined by this volume of  
 1353 IEEE Std 1003.1-200x shall behave as defined when called from or interrupted by a  
 1354 signal-catching function, with a single exception: when a signal interrupts an  
 1355 unsafe function and the signal-catching function calls an unsafe function, the  
 1356 behavior is undefined.

1357 When a signal is delivered to a thread, if the action of that signal specifies termination, stop, or  
 1358 continue, the entire process shall be terminated, stopped, or continued, respectively.

#### 1359 2.4.4 Signal Effects on Other Functions

1360 Signals affect the behavior of certain functions defined by this volume of IEEE Std 1003.1-200x if  
1361 delivered to a process while it is executing such a function. If the action of the signal is to  
1362 terminate the process, the process shall be terminated and the function shall not return. If the  
1363 action of the signal is to stop the process, the process shall stop until continued or terminated.  
1364 Generation of a SIGCONT signal for the process shall cause the process to be continued, and the  
1365 original function shall continue at the point the process was stopped. If the action of the signal is  
1366 to invoke a signal-catching function, the signal-catching function shall be invoked; in this case  
1367 the original function is said to be *interrupted* by the signal. If the signal-catching function  
1368 executes a **return** statement, the behavior of the interrupted function shall be as described  
1369 individually for that function, except as noted for unsafe functions. Signals that are ignored shall  
1370 not affect the behavior of any function; signals that are blocked shall not affect the behavior of  
1371 any function until they are unblocked and then delivered, except as specified for the *sigpending()*  
1372 and *sigwait()* functions.

#### 1373 2.5 Standard I/O Streams

1374 A stream is associated with an external file (which may be a physical device) by *opening* a file,  
1375 which may involve *creating* a new file. Creating an existing file causes its former contents to be  
1376 discarded if necessary. If a file can support positioning requests, (such as a disk file, as opposed  
1377 to a terminal), then a *file position indicator* associated with the stream is positioned at the start  
1378 (byte number 0) of the file, unless the file is opened with append mode, in which case it is  
1379 implementation-defined whether the file position indicator is initially positioned at the  
1380 beginning or end of the file. The file position indicator is maintained by subsequent reads,  
1381 writes, and positioning requests, to facilitate an orderly progression through the file. All input  
1382 takes place as if bytes were read by successive calls to *fgetc()*; all output takes place as if bytes  
1383 were written by successive calls to *fputc()*.

1384 When a stream is *unbuffered*, bytes are intended to appear from the source or at the destination  
1385 as soon as possible; otherwise, bytes may be accumulated and transmitted as a block. When a  
1386 stream is *fully buffered*, bytes are intended to be transmitted as a block when a buffer is filled.  
1387 When a stream is *line buffered*, bytes are intended to be transmitted as a block when a newline  
1388 byte is encountered. Furthermore, bytes are intended to be transmitted as a block when a buffer  
1389 is filled, when input is requested on an unbuffered stream, or when input is requested on a line-  
1390 buffered stream that requires the transmission of bytes. Support for these characteristics is  
1391 implementation-defined, and may be affected via *setbuf()* and *setvbuf()*.

1392 A file may be disassociated from a controlling stream by *closing* the file. Output streams are  
1393 flushed (any unwritten buffer contents are transmitted) before the stream is disassociated from  
1394 the file. The value of a pointer to a **FILE** object is unspecified after the associated file is closed  
1395 (including the standard streams).

1396 A file may be subsequently reopened, by the same or another program execution, and its  
1397 contents reclaimed or modified (if it can be repositioned at its start). If the *main()* function  
1398 returns to its original caller, or if the *exit()* function is called, all open files are closed (hence all  
1399 output streams are flushed) before program termination. Other paths to program termination,  
1400 such as calling *abort()*, need not close all files properly.

1401 The address of the **FILE** object used to control a stream may be significant; a copy of a **FILE**  
1402 object need not necessarily serve in place of the original.

1403 At program start-up, three streams are predefined and need not be opened explicitly: *standard*  
1404 *input* (for reading conventional input), *standard output* (for writing conventional output), and

1405 *standard error* (for writing diagnostic output). When opened, the standard error stream is not  
 1406 fully buffered; the standard input and standard output streams are fully buffered if and only if  
 1407 the stream can be determined not to refer to an interactive device.

### 1408 **2.5.1 Interaction of File Descriptors and Standard I/O Streams**

1409 cx This section describes the interaction of file descriptors and standard I/O streams. This  
 1410 functionality is an extension to the ISO C standard (and the rest of this section is not further CX  
 1411 shaded).

1412 An open file description may be accessed through a file descriptor, which is created using  
 1413 functions such as *open()* or *pipe()*, or through a stream, which is created using functions such as  
 1414 *fopen()* or *popen()*. Either a file descriptor or a stream is called a *handle* on the open file  
 1415 description to which it refers; an open file description may have several handles.

1416 Handles can be created or destroyed by explicit user action, without affecting the underlying  
 1417 open file description. Some of the ways to create them include *fcntl()*, *dup()*, *fdopen()*, *fileno()*,  
 1418 and *fork()*. They can be destroyed by at least *fclose()*, *close()*, and the *exec* functions.

1419 A file descriptor that is never used in an operation that could affect the file offset (for example,  
 1420 *read()*, *write()*, or *lseek()*) is not considered a handle for this discussion, but could give rise to one  
 1421 (for example, as a consequence of *fdopen()*, *dup()*, or *fork()*). This exception does not include the  
 1422 file descriptor underlying a stream, whether created with *fopen()* or *fdopen()*, so long as it is not  
 1423 used directly by the application to affect the file offset. The *read()* and *write()* functions  
 1424 implicitly affect the file offset; *lseek()* explicitly affects it.

1425 The result of function calls involving any one handle (the *active handle*) is defined elsewhere in  
 1426 this volume of IEEE Std 1003.1-200x, but if two or more handles are used, and any one of them is  
 1427 a stream, the application shall ensure that their actions are coordinated as described below. If  
 1428 this is not done, the result is undefined.

1429 A handle which is a stream is considered to be closed when either an *fclose()* or *freopen()* is  
 1430 executed on it (the result of *freopen()* is a new stream, which cannot be a handle on the same  
 1431 open file description as its previous value), or when the process owning that stream terminates  
 1432 with *exit()* or *abort()*. A file descriptor is closed by *close()*, *\_exit()*, or the *exec* functions when  
 1433 FD\_CLOEXEC is set on that file descriptor.

1434 For a handle to become the active handle, the application shall ensure that the actions below are  
 1435 performed between the last use of the handle (the current active handle) and the first use of the  
 1436 second handle (the future active handle). The second handle then becomes the active handle. All  
 1437 activity by the application affecting the file offset on the first handle shall be suspended until it  
 1438 again becomes the active file handle. (If a stream function has as an underlying function one that  
 1439 affects the file offset, the stream function shall be considered to affect the file offset.)

1440 The handles need not be in the same process for these rules to apply.

1441 Note that after a *fork()*, two handles exist where one existed before. The application shall ensure  
 1442 that, if both handles can ever be accessed, they are both in a state where the other could become  
 1443 the active handle first. The application shall prepare for a *fork()* exactly as if it were a change of  
 1444 active handle. (If the only action performed by one of the processes is one of the *exec* functions or  
 1445 *\_exit()* (not *exit()*), the handle is never accessed in that process.)

1446 For the first handle, the first applicable condition below applies. After the actions required  
 1447 below are taken, if the handle is still open, the application can close it.

- 1448 • If it is a file descriptor, no action is required.

- 1449       • If the only further action to be performed on any handle to this open file descriptor is to close  
1450       it, no action need be taken.
- 1451       • If it is a stream which is unbuffered, no action need be taken.
- 1452       • If it is a stream which is line buffered, and the last byte written to the stream was a newline  
1453       (that is, as if a:
- 1454             putc( '\n' )
- 1455       was the most recent operation on that stream), no action need be taken.
- 1456       • If it is a stream which is open for writing or appending (but not also open for reading), the  
1457       application shall either perform an *flush()*, or the stream shall be closed.
- 1458       • If the stream is open for reading and it is at the end of the file (*feof()* is true), no action need  
1459       be taken.
- 1460       • If the stream is open with a mode that allows reading and the underlying open file  
1461       description refers to a device that is capable of seeking, the application shall either perform  
1462       an *flush()*, or the stream shall be closed.
- 1463       Otherwise, the result is undefined.
- 1464       For the second handle:
- 1465       • If any previous active handle has been used by a function that explicitly changed the file  
1466       offset, except as required above for the first handle, the application shall perform an *lseek()* or  
1467       *fseek()* (as appropriate to the type of handle) to an appropriate location.
- 1468       If the active handle ceases to be accessible before the requirements on the first handle, above,  
1469       have been met, the state of the open file description becomes undefined. This might occur during  
1470       functions such as a *fork()* or *\_exit()*.
- 1471       The *exec* functions make inaccessible all streams that are open at the time they are called,  
1472       independent of which streams or file descriptors may be available to the new process image.
- 1473       When these rules are followed, regardless of the sequence of handles used, implementations  
1474       shall ensure that an application, even one consisting of several processes, shall yield correct  
1475       results: no data shall be lost or duplicated when writing, and all data shall be written in order,  
1476       except as requested by seeks. It is implementation-defined whether, and under what conditions,  
1477       all input is seen exactly once.
- 1478       If the rules above are not followed, the result is unspecified.
- 1479       Each function that operates on a stream is said to have zero or more *underlying functions*. This  
1480       means that the stream function shares certain traits with the underlying functions, but does not  
1481       require that there be any relation between the implementations of the stream function and its  
1482       underlying functions.

## 1483 2.5.2 Stream Orientation and Encoding Rules

1484       For conformance to the ISO/IEC 9899:1999 standard, the definition of a stream includes an  
1485       *orientation*. After a stream is associated with an external file, but before any operations are  
1486       performed on it, the stream is without orientation. Once a wide-character input/output function  
1487       has been applied to a stream without orientation, the stream shall become *wide-oriented*.  
1488       Similarly, once a byte input/output function has been applied to a stream without orientation,  
1489       the stream shall become *byte-oriented*. Only a call to the *freopen()* function or the *fwide()* function  
1490       can otherwise alter the orientation of a stream.

1491 A successful call to *freopen()* shall remove any orientation. The three predefined streams *standard*  
 1492 *input*, *standard output*, and *standard error* shall be unoriented at program start-up.

1493 Byte input/output functions cannot be applied to a wide-oriented stream, and wide-character  
 1494 input/output functions cannot be applied to a byte-oriented stream. The remaining stream  
 1495 operations shall not affect and shall not be affected by a stream's orientation, except for the  
 1496 following additional restrictions:

- 1497 • Binary wide-oriented streams have the file positioning restrictions ascribed to both text and  
 1498 binary streams.
- 1499 • For wide-oriented streams, after a successful call to a file-positioning function that leaves the  
 1500 file position indicator prior to the end-of-file, a wide-character output function can overwrite  
 1501 a partial character; any file contents beyond the byte(s) written are henceforth undefined.

1502 Each wide-oriented stream has an associated **mbstate\_t** object that stores the current parse state  
 1503 of the stream. A successful call to *fgetpos()* shall store a representation of the value of this  
 1504 **mbstate\_t** object as part of the value of the **fpos\_t** object. A later successful call to *fsetpos()* using  
 1505 the same stored **fpos\_t** value shall restore the value of the associated **mbstate\_t** object as well as  
 1506 the position within the controlled stream.

1507 Implementations that support multiple encoding rules associate an encoding rule with the  
 1508 stream. The encoding rule shall be determined by the setting of the *LC\_CTYPE* category in the  
 1509 current locale at the time when the stream becomes wide-oriented. If a wide-character  
 1510 input/output function is applied to a byte-oriented stream, the encoding rule used is undefined.  
 1511 As with the stream's orientation, the encoding rule associated with a stream cannot be changed  
 1512 once it has been set, except by a successful call to *freopen()* which clears the encoding rule and  
 1513 resets the orientation to unoriented.

1514 Although both text and binary wide-oriented streams are conceptually sequences of wide  
 1515 characters, the external file associated with a wide-oriented stream is a sequence of (possibly  
 1516 multi-byte) characters generalized as follows:

- 1517 • Multi-byte encodings within files may contain embedded null bytes (unlike multi-byte  
 1518 encodings valid for use internal to the program).
- 1519 • A file need not begin nor end in the initial shift state.

1520 Moreover, the encodings used for characters may differ among files. Both the nature and choice  
 1521 of such encodings are implementation-defined.

1522 The wide-character input functions read characters from the stream and convert them to wide  
 1523 characters as if they were read by successive calls to the *fgetwc()* function. Each conversion shall  
 1524 occur as if by a call to the *mbrtowc()* function, with the conversion state described by the stream's  
 1525 own **mbstate\_t** object, except the encoding rule associated with the stream is used instead of the  
 1526 encoding rule implied by the *LC\_CTYPE* category of the current locale.

1527 The wide-character output functions convert wide characters to (possibly multi-byte) characters  
 1528 and write them to the stream as if they were written by successive calls to the *fputwc()* function.  
 1529 Each conversion shall occur as if by a call to the *wcrtomb()* function, with the conversion state  
 1530 described by the stream's own **mbstate\_t** object, except the encoding rule associated with the  
 1531 stream is used instead of the encoding rule implied by the *LC\_CTYPE* category of the current  
 1532 locale.

1533 An *encoding error* shall occur if the character sequence presented to the underlying *mbrtowc()*  
 1534 function does not form a valid (generalized) character, or if the code value passed to the  
 1535 underlying *wcrtomb()* function does not correspond to a valid (generalized) character. The  
 1536 wide-character input/output functions and the byte input/output functions store the value of

1537 the macro [EILSEQ] in *errno* if and only if an encoding error occurs.

## 1538 **2.6 STREAMS**

1539 XSR STREAMS functionality is provided on implementations supporting the XSI STREAMS Option  
1540 Group. This functionality is dependent on support of the XSI STREAMS option (and the rest of  
1541 this section is not further shaded for this option).

1542 STREAMS provides a uniform mechanism for implementing networking services and other  
1543 character-based I/O. The STREAMS function provides direct access to protocol modules.  
1544 STREAMS modules are unspecified objects. Access to STREAMS modules is provided by  
1545 interfaces in IEEE Std 1003.1-200x. Creation of STREAMS modules is outside the scope of  
1546 IEEE Std 1003.1-200x.

1547 A STREAM is typically a full-duplex connection between a process and an open device or  
1548 pseudo-device. However, since pipes may be STREAMS-based, a STREAM can be a full-duplex  
1549 connection between two processes. The STREAM itself exists entirely within the implementation  
1550 and provides a general character I/O function for processes. It optionally includes one or more  
1551 intermediate processing modules that are interposed between the process end of the STREAM  
1552 (STREAM head) and a device driver at the end of the STREAM (STREAM end).

1553 STREAMS I/O is based on messages. There are three types of message: |

- 1554 • *Data messages* containing actual data for input or output
- 1555 • *Control data* containing instructions for the STREAMS modules and underlying  
1556 implementation
- 1557 • Other messages, which include file descriptors

1558 The interface between the STREAM and the rest of the implementation is provided by a set of |  
1559 functions at the STREAM head. When a process calls *write()*, *writew()*, *putmsg()*, *putpmsg()*, or |  
1560 *ioctl()*, messages are sent down the STREAM, and *read()*, *readv()*, *getmsg()*, or *getpmsg()* accepts |  
1561 data from the STREAM and passes it to a process. Data intended for the device at the  
1562 downstream end of the STREAM is packaged into messages and sent downstream, while data  
1563 and signals from the device are composed into messages by the device driver and sent upstream  
1564 to the STREAM head.

1565 When a STREAMS-based device is opened, a STREAM shall be created that contains the |  
1566 STREAM head and the STREAM end (driver). If pipes are STREAMS-based in an |  
1567 implementation, when a pipe is created, two STREAMS shall be created, each containing a |  
1568 STREAM head. Other modules are added to the STREAM using *ioctl()*. New modules are |  
1569 “pushed” onto the STREAM one at a time in last-in, first-out (LIFO) style, as though the |  
1570 STREAM was a push-down stack.

### 1571 **Priority**

1572 Message types are classified according to their queuing priority and may be *normal* (non-  
1573 priority), *priority*, or *high-priority* messages. A message belongs to a particular priority band that  
1574 determines its ordering when placed on a queue. Normal messages have a priority band of 0 and |  
1575 shall always be placed at the end of the queue following all other messages in the queue. High-  
1576 priority messages are always placed at the head of a queue, but shall be discarded if there is |  
1577 already a high-priority message in the queue. Their priority band shall be ignored; they are |  
1578 high-priority by virtue of their type. Priority messages have a priority band greater than 0. |  
1579 Priority messages are always placed after any messages of the same or higher priority. High-  
1580 priority and priority messages are used to send control and data information outside the normal |

1581 flow of control. By convention, high-priority messages shall not be affected by flow control. |  
 1582 Normal and priority messages have separate flow controls. |

### 1583 **Message Parts**

1584 A process may access STREAMS messages that contain a data part, control part, or both. The  
 1585 data part is that information which is transmitted over the communication medium and the  
 1586 control information is used by the local STREAMS modules. The other types of messages are  
 1587 used between modules and are not accessible to processes. Messages containing only a data part  
 1588 are accessible via *putmsg()*, *putpmsg()*, *getmsg()*, *getpmsg()*, *read()*, *readv()*, *write()*, or *writv()*. |  
 1589 Messages containing a control part with or without a data part are accessible via calls to  
 1590 *putmsg()*, *putpmsg()*, *getmsg()*, or *getpmsg()*.

### 1591 **2.6.1 Accessing STREAMS**

1592 A process accesses STREAMS-based files using the standard functions *close()*, *ioctl()*, *getmsg()*,  
 1593 *getpmsg()*, *open()*, *pipe()*, *poll()*, *putmsg()*, *putpmsg()*, *read()*, or *write()*. Refer to the applicable  
 1594 function definitions for general properties and errors.

1595 Calls to *ioctl()* shall perform control functions on the STREAM associated with the file descriptor |  
 1596 *fildev*. The control functions may be performed by the STREAM head, a STREAMS module, or |  
 1597 the STREAMS driver for the STREAM. |

1598 STREAMS modules and drivers can detect errors, sending an error message to the STREAM |  
 1599 head, thus causing subsequent functions to fail and set *errno* to the value specified in the |  
 1600 message. In addition, STREAMS modules and drivers can elect to fail a particular *ioctl()* request |  
 1601 alone by sending a negative acknowledgement message to the STREAM head. This shall cause |  
 1602 just the pending *ioctl()* request to fail and set *errno* to the value specified in the message. |

## 1603 **2.7 XSI Interprocess Communication**

1604 XSI This section describes extensions to support interprocess communication. This functionality is  
 1605 dependent on support of the XSI Extension (and the rest of this section is not further shaded for  
 1606 this option).

1607 The following message passing, semaphore, and shared memory services form an XSI  
 1608 interprocess communication facility. Certain aspects of their operation are common, and are  
 1609 described below.

1610

1611

1612

1613

1614

1615

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>	

1616 Another interprocess communication facility is provided by functions in the Realtime Option  
 1617 Group; see Section 2.8 (on page 491).

1618 **2.7.1 IPC General Description**

1619 Each individual shared memory segment, message queue, and semaphore set shall be identified |  
 1620 by a unique positive integer, called, respectively, a shared memory identifier, *shmid*, a |  
 1621 semaphore identifier, *semid*, and a message queue identifier, *msgid*. The identifiers shall be |  
 1622 returned by calls to *shmget()*, *semget()*, and *msgget()*, respectively. |

1623 Associated with each identifier is a data structure which contains data related to the operations  
 1624 which may be or may have been performed; see the Base Definitions volume of  
 1625 IEEE Std 1003.1-200x, <*sys/shm.h*>, <*sys/sem.h*>, and <*sys/msg.h*> for their descriptions.

1626 Each of the data structures contains both ownership information and an **ipc\_perm** structure (see  
 1627 the Base Definitions volume of IEEE Std 1003.1-200x, <*sys/ipc.h*>) which are used in conjunction  
 1628 to determine whether or not read/write (read/alter for semaphores) permissions should be  
 1629 granted to processes using the IPC facilities. The *mode* member of the **ipc\_perm** structure acts as  
 1630 a bit field which determines the permissions.

1631 The values of the bits are given below in octal notation.

1632  
 1633

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.

1634  
 1635  
 1636  
 1637  
 1638  
 1639

1640 The name of the **ipc\_perm** structure is *shm\_perm*, *sem\_perm*, or *msg\_perm*, depending on which  
 1641 service is being used. In each case, read and write/alter permissions shall be granted to a process |  
 1642 if one or more of the following are true ("xxx" is replaced by *shm*, *sem*, or *msg*, as appropriate): |

- 1643 • The process has appropriate privileges.
- 1644 • The effective user ID of the process matches *xxx\_perm.cuid* or *xxx\_perm.uid* in the data  
 1645 structure associated with the IPC identifier, and the appropriate bit of the *user* field in  
 1646 *xxx\_perm.mode* is set.
- 1647 • The effective user ID of the process does not match *xxx\_perm.cuid* or *xxx\_perm.uid* but the  
 1648 effective group ID of the process matches *xxx\_perm.cgid* or *xxx\_perm.gid* in the data structure  
 1649 associated with the IPC identifier, and the appropriate bit of the *group* field in *xxx\_perm.mode*  
 1650 is set.
- 1651 • The effective user ID of the process does not match *xxx\_perm.cuid* or *xxx\_perm.uid* and the  
 1652 effective group ID of the process does not match *xxx\_perm.cgid* or *xxx\_perm.gid* in the data  
 1653 structure associated with the IPC identifier, but the appropriate bit of the *other* field in  
 1654 *xxx\_perm.mode* is set.

1655 Otherwise, the permission shall be denied. |



## 1656 2.8 Realtime

1657 This section defines functions to support the source portability of applications with realtime  
1658 requirements. The presence of many of these functions is dependent on support for  
1659 implementation options described in the text.

1660 The specific functional areas included in this section and their scope include the following. Full  
1661 definitions of these terms can be found in the Base Definitions volume of IEEE Std 1003.1-200x,  
1662 Chapter 3, Definitions.

- 1663 • Semaphores
- 1664 • Process Memory Locking
- 1665 • Memory Mapped Files and Shared Memory Objects
- 1666 • Priority Scheduling
- 1667 • Realtime Signal Extension
- 1668 • Timers
- 1669 • Interprocess Communication
- 1670 • Synchronized Input and Output
- 1671 • Asynchronous Input and Output

1672 All the realtime functions defined in this volume of IEEE Std 1003.1-200x are portable, although  
1673 some of the numeric parameters used by an implementation may have hardware dependencies.

### 1674 2.8.1 Realtime Signals

1675 RTS Realtime signal generation and delivery is dependent on support for the Realtime Signals  
1676 Extension option.

1677 See Section 2.4.2 (on page 479).

### 1678 2.8.2 Asynchronous I/O

1679 AIO The functionality described in this section is dependent on support of the Asynchronous Input  
1680 and Output option (and the rest of this section is not further shaded for this option).

1681 An asynchronous I/O control block structure **aiocb** is used in many asynchronous I/O  
1682 functions. It is defined in the Base Definitions volume of IEEE Std 1003.1-200x, <**aio.h**> and has  
1683 at least the following members:

1684	Member Type	Member Name	Description
1685	<b>int</b>	<i>aio_fildes</i>	File descriptor.
1686	<b>off_t</b>	<i>aio_offset</i>	File offset.
1687	<b>volatile void*</b>	<i>aio_buf</i>	Location of buffer.
1688	<b>size_t</b>	<i>aio_nbytes</i>	Length of transfer.
1689	<b>int</b>	<i>aio_reqprio</i>	Request priority offset.
1690	<b>struct sigevent</b>	<i>aio_sigevent</i>	Signal number and value.
1691	<b>int</b>	<i>aio_lio_opcode</i>	Operation to be performed.

1692 The *aio\_fildes* element is the file descriptor on which the asynchronous operation is performed.

1693 If **O\_APPEND** is not set for the file descriptor *aio\_fildes* and if *aio\_fildes* is associated with a  
1694 device that is capable of seeking, then the requested operation takes place at the absolute  
1695 position in the file as given by *aio\_offset*, as if *lseek()* were called immediately prior to the

1696 operation with an *offset* argument equal to *aio\_offset* and a *whence* argument equal to `SEEK_SET`.  
1697 If `O_APPEND` is set for the file descriptor, or if *aio\_fildes* is associated with a device that is  
1698 incapable of seeking, write operations append to the file in the same order as the calls were  
1699 made, with the following exception: under implementation-defined circumstances, such as  
1700 operation on a multi-processor or when requests of differing priorities are submitted at the same  
1701 time, the ordering restriction may be relaxed. Since there is no way for a strictly conforming  
1702 application to determine whether this relaxation applies, all strictly conforming applications  
1703 which rely on ordering of output shall be written in such a way that they will operate correctly if  
1704 the relaxation applies. After a successful call to enqueue an asynchronous I/O operation, the  
1705 value of the file offset for the file is unspecified. The *aio\_nbytes* and *aio\_buf* elements are the same  
1706 as the *nbyte* and *buf* arguments defined by `read()` and `write()`, respectively.

1707 If `_POSIX_PRIORITIZED_IO` and `_POSIX_PRIORITY_SCHEDULING` are defined, then  
1708 asynchronous I/O is queued in priority order, with the priority of each asynchronous operation  
1709 based on the current scheduling priority of the calling process. The *aio\_reqprio* member can be  
1710 used to lower (but not raise) the asynchronous I/O operation priority and is within the range  
1711 zero through `{AIO_PRIO_DELTA_MAX}`, inclusive. Unless both `_POSIX_PRIORITIZED_IO` and  
1712 `_POSIX_PRIORITY_SCHEDULING` are defined, the order of processing asynchronous I/O  
1713 requests is unspecified. When both `_POSIX_PRIORITIZED_IO` and  
1714 `_POSIX_PRIORITY_SCHEDULING` are defined, the order of processing of requests submitted  
1715 by processes whose schedulers are not `SCHED_FIFO`, `SCHED_RR`, or `SCHED_SPORADIC` is  
1716 unspecified. The priority of an asynchronous request is computed as (process scheduling  
1717 priority) minus *aio\_reqprio*. The priority assigned to each asynchronous I/O request is an  
1718 indication of the desired order of execution of the request relative to other asynchronous I/O  
1719 requests for this file. If `_POSIX_PRIORITIZED_IO` is defined, requests issued with the same  
1720 priority to a character special file are processed by the underlying device in FIFO order; the order  
1721 of processing of requests of the same priority issued to files that are not character special files is  
1722 unspecified. Numerically higher priority values indicate requests of higher priority. The value of  
1723 *aio\_reqprio* has no effect on process scheduling priority. When prioritized asynchronous I/O  
1724 requests to the same file are blocked waiting for a resource required for that I/O operation, the  
1725 higher-priority I/O requests shall be granted the resource before lower-priority I/O requests are  
1726 granted the resource. The relative priority of asynchronous I/O and synchronous I/O is  
1727 implementation-defined. If `_POSIX_PRIORITIZED_IO` is defined, the implementation shall  
1728 define for which files I/O prioritization is supported.

1729 The *aio\_sigevent* determines how the calling process shall be notified upon I/O completion, as  
1730 specified in Section 2.4.1 (on page 478). If *aio\_sigevent.sigev\_notify* is `SIGEV_NONE`, then no  
1731 signal shall be posted upon I/O completion, but the error status for the operation and the return  
1732 status for the operation shall be set appropriately.

1733 The *aio\_lio\_opcode* field is used only by the `lio_listio()` call. The `lio_listio()` call allows multiple  
1734 asynchronous I/O operations to be submitted at a single time. The function takes as an  
1735 argument an array of pointers to **aiocb** structures. Each **aiocb** structure indicates the operation to  
1736 be performed (read or write) via the *aio\_lio\_opcode* field.

1737 The address of the **aiocb** structure is used as a handle for retrieving the error status and return  
1738 status of the asynchronous operation while it is in progress.

1739 The **aiocb** structure and the data buffers associated with the asynchronous I/O operation are  
1740 being used by the system for asynchronous I/O while, and only while, the error status of the  
1741 asynchronous operation is equal to `[EINPROGRESS]`. Applications shall not modify the **aiocb**  
1742 structure while the structure is being used by the system for asynchronous I/O.

1743 The return status of the asynchronous operation is the number of bytes transferred by the I/O  
1744 operation. If the error status is set to indicate an error completion, then the return status is set to

1745 the return value that the corresponding *read()*, *write()*, or *fsync()* call would have returned.  
 1746 When the error status is not equal to [EINPROGRESS], the return status shall reflect the return  
 1747 status of the corresponding synchronous operation.

### 1748 **2.8.3 Memory Management**

#### 1749 *2.8.3.1 Memory Locking*

1750 ML The functionality described in this section is dependent on support of the Process Memory  
 1751 Locking option (and the rest of this section is not further shaded for this option).

1752 Range memory locking operations are defined in terms of pages. Implementations may restrict  
 1753 the size and alignment of range lockings to be on page-size boundaries. The page size, in bytes,  
 1754 is the value of the configurable system variable {PAGESIZE}. If an implementation has no  
 1755 restrictions on size or alignment, it may specify a 1-byte page size.

1756 Memory locking guarantees the residence of portions of the address space. It is  
 1757 implementation-defined whether locking memory guarantees fixed translation between virtual  
 1758 addresses (as seen by the process) and physical addresses. Per-process memory locks are not  
 1759 inherited across a *fork()*, and all memory locks owned by a process are unlocked upon *exec* or  
 1760 process termination. Unmapping of an address range removes any memory locks established on  
 1761 that address range by this process.

#### 1762 *2.8.3.2 Memory Mapped Files*

1763 MF The functionality described in this section is dependent on support of the Memory Mapped Files  
 1764 option (and the rest of this section is not further shaded for this option).

1765 Range memory mapping operations are defined in terms of pages. Implementations may  
 1766 restrict the size and alignment of range mappings to be on page-size boundaries. The page size,  
 1767 in bytes, is the value of the configurable system variable {PAGESIZE}. If an implementation has  
 1768 no restrictions on size or alignment, it may specify a 1-byte page size.

1769 Memory mapped files provide a mechanism that allows a process to access files by directly  
 1770 incorporating file data into its address space. Once a file is mapped into a process address space,  
 1771 the data can be manipulated as memory. If more than one process maps a file, its contents are  
 1772 shared among them. If the mappings allow shared write access, then data written into the  
 1773 memory object through the address space of one process appears in the address spaces of all  
 1774 processes that similarly map the same portion of the memory object.

1775 SHM Shared memory objects are named regions of storage that may be independent of the file system  
 1776 and can be mapped into the address space of one or more processes to allow them to share the  
 1777 associated memory.

1778 SHM An *unlink()* of a file or *shm\_unlink()* of a shared memory object, while causing the removal of the  
 1779 name, does not unmap any mappings established for the object. Once the name has been  
 1780 removed, the contents of the memory object are preserved as long as it is referenced. The  
 1781 memory object remains referenced as long as a process has the memory object open or has some  
 1782 area of the memory object mapped.

### 1783 2.8.3.3 *Memory Protection*

1784 MPR MF The functionality described in this section is dependent on support of the Memory Protection  
1785 and Memory Mapped Files option (and the rest of this section is not further shaded for these  
1786 options).

1787 When an object is mapped, various application accesses to the mapped region may result in  
1788 signals. In this context, SIGBUS is used to indicate an error using the mapped object, and  
1789 SIGSEGV is used to indicate a protection violation or misuse of an address:

- 1790 • A mapping may be restricted to disallow some types of access.
- 1791 • Write attempts to memory that was mapped without write access, or any access to memory  
1792 mapped PROT\_NONE, shall result in a SIGSEGV signal.
- 1793 • References to unmapped addresses shall result in a SIGSEGV signal.
- 1794 • Reference to whole pages within the mapping, but beyond the current length of the object,  
1795 shall result in a SIGBUS signal.
- 1796 • The size of the object is unaffected by access beyond the end of the object (even if a SIGBUS is  
1797 not generated).

### 1798 2.8.3.4 *Typed Memory Objects*

1799 TYM The functionality described in this section is dependent on support of the Typed Memory  
1800 Objects option (and the rest of this section is not further shaded for this option).

1801 Implementations may support the Typed Memory Objects option without supporting the  
1802 Memory Mapped Files option or the Shared Memory Objects option. Typed memory objects are  
1803 implementation-configurable named storage pools accessible from one or more processors in a  
1804 system, each via one or more ports, such as backplane buses, LANs, I/O channels, and so on.  
1805 Each valid combination of a storage pool and a port is identified through a name that is defined  
1806 at system configuration time, in an implementation-defined manner; the name may be  
1807 independent of the file system. Using this name, a typed memory object can be opened and  
1808 mapped into process address space. For a given storage pool and port, it is necessary to support  
1809 both dynamic allocation from the pool as well as mapping at an application-supplied offset  
1810 within the pool; when dynamic allocation has been performed, subsequent deallocation must be  
1811 supported. Lastly, accessing typed memory objects from different ports requires a method for  
1812 obtaining the offset and length of contiguous storage of a region of typed memory (dynamically  
1813 allocated or not); this allows typed memory to be shared among processes and/or processors  
1814 while being accessed from the desired port.

## 1815 2.8.4 **Process Scheduling**

1816 PS The functionality described in this section is dependent on support of the Process Scheduling  
1817 option (and the rest of this section is not further shaded for this option).

### 1818 **Scheduling Policies**

1819 The scheduling semantics described in this volume of IEEE Std 1003.1-200x are defined in terms  
1820 of a conceptual model that contains a set of thread lists. No implementation structures are  
1821 necessarily implied by the use of this conceptual model. It is assumed that no time elapses  
1822 during operations described using this model, and therefore no simultaneous operations are  
1823 possible. This model discusses only processor scheduling for runnable threads, but it should be  
1824 noted that greatly enhanced predictability of realtime applications results if the sequencing of  
1825 other resources takes processor scheduling policy into account.

1826 There is, conceptually, one thread list for each priority. A runnable thread will be on the thread |  
 1827 list for that thread's priority. Multiple scheduling policies shall be provided. Each non-empty |  
 1828 thread list is ordered, contains a head as one end of its order, and a tail as the other. The purpose  
 1829 of a scheduling policy is to define the allowable operations on this set of lists (for example,  
 1830 moving threads between and within lists).

1831 Each process shall be controlled by an associated scheduling policy and priority. These  
 1832 parameters may be specified by explicit application execution of the *sched\_setscheduler()* or  
 1833 *sched\_setparam()* functions.

1834 Each thread shall be controlled by an associated scheduling policy and priority. These  
 1835 parameters may be specified by explicit application execution of the *pthread\_setschedparam()*  
 1836 function.

1837 Associated with each policy is a priority range. Each policy definition shall specify the minimum  
 1838 priority range for that policy. The priority ranges for each policy may but need not overlap the  
 1839 priority ranges of other policies.

1840 A conforming implementation shall select the thread that is defined as being at the head of the  
 1841 highest priority non-empty thread list to become a running thread, regardless of its associated  
 1842 policy. This thread is then removed from its thread list.

1843 Four scheduling policies are specifically required. Other implementation-defined scheduling  
 1844 policies may be defined. The following symbols are defined in the Base Definitions volume of  
 1845 IEEE Std 1003.1-200x, <**sched.h**>:

1846 **SCHED\_FIFO** First in, first out (FIFO) scheduling policy.

1847 **SCHED\_RR** Round robin scheduling policy.

1848 ss **SCHED\_SPORADIC** Sporadic server scheduling policy.

1849 **SCHED\_OTHER** Another scheduling policy.

1850 The values of these symbols shall be distinct.

### 1851 **SCHED\_FIFO**

1852 Conforming implementations shall include a scheduling policy called the FIFO scheduling  
 1853 policy.

1854 Threads scheduled under this policy are chosen from a thread list that is ordered by the time its  
 1855 threads have been on the list without being executed; generally, the head of the list is the thread  
 1856 that has been on the list the longest time, and the tail is the thread that has been on the list the  
 1857 shortest time.

1858 Under the **SCHED\_FIFO** policy, the modification of the definitional thread lists is as follows:

- 1859 1. When a running thread becomes a preempted thread, it becomes the head of the thread list  
 1860 for its priority.
- 1861 2. When a blocked thread becomes a runnable thread, it becomes the tail of the thread list for  
 1862 its priority.
- 1863 3. When a running thread calls the *sched\_setscheduler()* function, the process specified in the  
 1864 function call is modified to the specified policy and the priority specified by the *param*  
 1865 argument.
- 1866 4. When a running thread calls the *sched\_setparam()* function, the priority of the process  
 1867 specified in the function call is modified to the priority specified by the *param* argument.

- 1868 5. When a running thread calls the *pthread\_setschedparam()* function, the thread specified in  
 1869 the function call is modified to the specified policy and the priority specified by the *param*  
 1870 argument.
- 1871 6. When a running thread calls the *pthread\_setschedprio()* function, the thread specified in the  
 1872 function call is modified to the priority specified by the *prio* argument.
- 1873 7. If a thread whose policy or priority has been modified other than by *pthread\_setschedprio()*  
 1874 is a running thread or is runnable, it then becomes the tail of the thread list for its new  
 1875 priority.
- 1876 8. If a thread whose policy or priority has been modified by *pthread\_setschedprio()* is a  
 1877 running thread or is runnable, the effect on its position in the thread list depends on the  
 1878 direction of the modification, as follows:
- 1879 a. If the priority is raised, the thread becomes the tail of the thread list.
- 1880 b. If the priority is unchanged, the thread does not change position in the thread list.
- 1881 c. If the priority is lowered, the thread becomes the head of the thread list.
- 1882 9. When a running thread issues the *sched\_yield()* function, the thread becomes the tail of the  
 1883 thread list for its priority.
- 1884 10. At no other time is the position of a thread with this scheduling policy within the thread  
 1885 lists affected.

1886 For this policy, valid priorities shall be within the range returned by the *sched\_get\_priority\_max()*  
 1887 and *sched\_get\_priority\_min()* functions when SCHED\_FIFO is provided as the parameter.  
 1888 Conforming implementations shall provide a priority range of at least 32 priorities for this  
 1889 policy.

## 1890 SCHED\_RR

1891 Conforming implementations shall include a scheduling policy called the *round robin* scheduling  
 1892 policy. This policy shall be identical to the SCHED\_FIFO policy with the additional condition  
 1893 that when the implementation detects that a running thread has been executing as a running  
 1894 thread for a time period of the length returned by the *sched\_rr\_get\_interval()* function or longer,  
 1895 the thread shall become the tail of its thread list and the head of that thread list shall be removed  
 1896 and made a running thread.

1897 The effect of this policy is to ensure that if there are multiple SCHED\_RR threads at the same  
 1898 priority, one of them does not monopolize the processor. An application should not rely only on  
 1899 the use of SCHED\_RR to ensure application progress among multiple threads if the application  
 1900 includes threads using the SCHED\_FIFO policy at the same or higher priority levels or  
 1901 SCHED\_RR threads at a higher priority level.

1902 A thread under this policy that is preempted and subsequently resumes execution as a running  
 1903 thread completes the unexpired portion of its round robin interval time period.

1904 For this policy, valid priorities shall be within the range returned by the *sched\_get\_priority\_max()*  
 1905 and *sched\_get\_priority\_min()* functions when SCHED\_RR is provided as the parameter.  
 1906 Conforming implementations shall provide a priority range of at least 32 priorities for this  
 1907 policy.

1908 **SCHED\_SPORADIC**

1909 SS|TSP The functionality described in this section is dependent on support of the Process Sporadic  
 1910 Server or Thread Sporadic Server options (and the rest of this section is not further shaded for  
 1911 these options).

1912 If `_POSIX_SPORADIC_SERVER` or `_POSIX_THREAD_SPORADIC_SERVER` is defined, the  
 1913 implementation shall include a scheduling policy identified by the value `SCHED_SPORADIC`.

1914 The sporadic server policy is based primarily on two parameters: the *replenishment period* and the  
 1915 *available execution capacity*. The replenishment period is given by the *sched\_ss\_repl\_period*  
 1916 member of the **sched\_param** structure. The available execution capacity is initialized to the  
 1917 value given by the *sched\_ss\_init\_budget* member of the same parameter. The sporadic server  
 1918 policy is identical to the `SCHED_FIFO` policy with some additional conditions that cause the  
 1919 thread's assigned priority to be switched between the values specified by the *sched\_priority* and  
 1920 *sched\_ss\_low\_priority* members of the **sched\_param** structure.

1921 The priority assigned to a thread using the sporadic server scheduling policy is determined in  
 1922 the following manner: if the available execution capacity is greater than zero and the number of  
 1923 pending replenishment operations is strictly less than *sched\_ss\_max\_repl*, the thread is assigned  
 1924 the priority specified by *sched\_priority*; otherwise, the assigned priority shall be  
 1925 *sched\_ss\_low\_priority*. If the value of *sched\_priority* is less than or equal to the value of  
 1926 *sched\_ss\_low\_priority*, the results are undefined. When active, the thread shall belong to the  
 1927 thread list corresponding to its assigned priority level, according to the mentioned priority  
 1928 assignment. The modification of the available execution capacity and, consequently of the  
 1929 assigned priority, is done as follows:

- 1930 1. When the thread at the head of the *sched\_priority* list becomes a running thread, its  
 1931 execution time shall be limited to at most its available execution capacity, plus the  
 1932 resolution of the execution time clock used for this scheduling policy. This resolution shall  
 1933 be implementation-defined.
- 1934 2. Each time the thread is inserted at the tail of the list associated with *sched\_priority*—  
 1935 because as a blocked thread it became runnable with priority *sched\_priority* or because a  
 1936 replenishment operation was performed—the time at which this operation is done is  
 1937 posted as the *activation\_time*.
- 1938 3. When the running thread with assigned priority equal to *sched\_priority* becomes a  
 1939 preempted thread, it becomes the head of the thread list for its priority, and the execution  
 1940 time consumed is subtracted from the available execution capacity. If the available  
 1941 execution capacity would become negative by this operation, it shall be set to zero.
- 1942 4. When the running thread with assigned priority equal to *sched\_priority* becomes a blocked  
 1943 thread, the execution time consumed is subtracted from the available execution capacity,  
 1944 and a replenishment operation is scheduled, as described in 6 and 7. If the available  
 1945 execution capacity would become negative by this operation, it shall be set to zero.
- 1946 5. When the running thread with assigned priority equal to *sched\_priority* reaches the limit  
 1947 imposed on its execution time, it becomes the tail of the thread list for  
 1948 *sched\_ss\_low\_priority*, the execution time consumed is subtracted from the available  
 1949 execution capacity (which becomes zero), and a replenishment operation is scheduled, as  
 1950 described in 6 and 7.
- 1951 6. Each time a replenishment operation is scheduled, the amount of execution capacity to be  
 1952 replenished, *replenish\_amount*, is set equal to the execution time consumed by the thread  
 1953 since the *activation\_time*. The replenishment is scheduled to occur at *activation\_time* plus  
 1954 *sched\_ss\_repl\_period*. If the scheduled time obtained is before the current time, the

1955 replenishment operation is carried out immediately. Several replenishment operations may  
 1956 be pending at the same time, each of which will be serviced at its respective scheduled  
 1957 time. With the above rules, the number of replenishment operations simultaneously  
 1958 pending for a given thread that is scheduled under the sporadic server policy shall not be  
 1959 greater than *sched\_ss\_max\_repl*.

1960 7. A replenishment operation consists of adding the corresponding *replenish\_amount* to the  
 1961 available execution capacity at the scheduled time. If, as a consequence of this operation,  
 1962 the execution capacity would become larger than *sched\_ss\_initial\_budget*, it shall be  
 1963 rounded down to a value equal to *sched\_ss\_initial\_budget*. Additionally, if the thread was  
 1964 runnable or running, and had assigned priority equal to *sched\_ss\_low\_priority*, then it  
 1965 becomes the tail of the thread list for *sched\_priority*.

1966 Execution time is defined in Section 2.2.2 (on page 464).

1967 For this policy, changing the value of a CPU-time clock via *clock\_settime()* shall have no effect on  
 1968 its behavior.

1969 For this policy, valid priorities shall be within the range returned by the *sched\_get\_priority\_min()*  
 1970 and *sched\_get\_priority\_max()* functions when SCHED\_SPORADIC is provided as the parameter.  
 1971 Conforming implementations shall provide a priority range of at least 32 distinct priorities for  
 1972 this policy.

1973 **SCHED\_OTHER**

1974 Conforming implementations shall include one scheduling policy identified as SCHED\_OTHER  
 1975 (which may execute identically with either the FIFO or round robin scheduling policy). The  
 1976 effect of scheduling threads with the SCHED\_OTHER policy in a system in which other threads  
 1977 ss are executing under SCHED\_FIFO, SCHED\_RR, or SCHED\_SPORADIC is implementation-  
 1978 defined.

1979 This policy is defined to allow strictly conforming applications to be able to indicate in a  
 1980 portable manner that they no longer need a realtime scheduling policy.

1981 For threads executing under this policy, the implementation shall use only priorities within the  
 1982 range returned by the *sched\_get\_priority\_max()* and *sched\_get\_priority\_min()* functions when  
 1983 SCHED\_OTHER is provided as the parameter.

1984 **2.8.5 Clocks and Timers**

1985 TMR The functionality described in this section is dependent on support of the Timers option (and the  
 1986 rest of this section is not further shaded for this option).

1987 The <time.h> header defines the types and manifest constants used by the timing facility.

1988 **Time Value Specification Structures**

1989 Many of the timing facility functions accept or return time value specifications. A time value  
 1990 structure **timespec** specifies a single time value and includes at least the following members:

1991

1992

1993

1994

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

1995 The *tv\_nsec* member is only valid if greater than or equal to zero, and less than the number of  
 1996 nanoseconds in a second (1,000 million). The time interval described by this structure is (*tv\_sec* \*  
 1997 10<sup>9</sup> + *tv\_nsec*) nanoseconds.



1998 A time value structure **itimerspec** specifies an initial timer value and a repetition interval for use  
 1999 by the per-process timer functions. This structure includes at least the following members:

2000  
 2001  
 2002  
 2003

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

2004 If the value described by *it\_value* is non-zero, it indicates the time to or time of the next timer  
 2005 expiration (for relative and absolute timer values, respectively). If the value described by *it\_value*  
 2006 is zero, the timer shall be disarmed.

2007 If the value described by *it\_interval* is non-zero, it specifies an interval which shall be used in  
 2008 reloading the timer when it expires; that is, a periodic timer is specified. If the value described by  
 2009 *it\_interval* is zero, the timer is disarmed after its next expiration; that is, a one-shot timer is  
 2010 specified.

2011 **Timer Event Notification Control Block**

2012 RTS Per-process timers may be created that notify the process of timer expirations by queuing a  
 2013 realtime extended signal. The **sigevent** structure, defined in the Base Definitions volume of  
 2014 IEEE Std 1003.1-200x, <**signal.h**>, is used in creating such a timer. The **sigevent** structure  
 2015 contains the signal number and an application-specific data value which shall be used when  
 2016 notifying the calling process of timer expiration events.

2017 **Manifest Constants**

2018 The following constants are defined in the Base Definitions volume of IEEE Std 1003.1-200x,  
 2019 <**time.h**>:

2020 **CLOCK\_REALTIME** The identifier for the system-wide realtime clock.

2021 **TIMER\_ABSTIME** Flag indicating time is absolute with respect to the clock associated  
 2022 with a timer.

2023 MON **CLOCK\_MONOTONIC** The identifier for the system-wide monotonic clock, which is defined  
 2024 as a clock whose value cannot be set via *clock\_settime()* and which  
 2025 cannot have backward clock jumps. The maximum possible clock  
 2026 jump is implementation-defined.

2027 MON The maximum allowable resolution for **CLOCK\_REALTIME** and **CLOCK\_MONOTONIC** clocks  
 2028 and all time services based on these clocks is represented by **{\_POSIX\_CLOCKRES\_MIN}** and  
 2029 shall be defined as 20ms (1/50 of a second). Implementations may support smaller values of  
 2030 resolution for these clocks to provide finer granularity time bases. The actual resolution  
 2031 supported by an implementation for a specific clock is obtained using the *clock\_getres()* function.  
 2032 If the actual resolution supported for a time service based on one of these clocks differs from the  
 2033 resolution supported for that clock, the implementation shall document this difference.

2034 MON The minimum allowable maximum value for **CLOCK\_REALTIME** and **CLOCK\_MONOTONIC**  
 2035 clocks and all absolute time services based on them is the same as that defined by the ISO C  
 2036 standard for the **time\_t** type. If the maximum value supported by a time service based on one of  
 2037 these clocks differs from the maximum value supported by that clock, the implementation shall  
 2038 document this difference.

2039 **Execution Time Monitoring**

2040 CPT If `_POSIX_CPUTIME` is defined, process CPU-time clocks shall be supported in addition to the  
2041 clocks described in **Manifest Constants** (on page 499).

2042 TCT If `_POSIX_THREAD_CPUTIME` is defined, thread CPU-time clocks shall be supported.

2043 CPT|TCT CPU-time clocks measure execution or CPU time, which is defined in the Base Definitions  
2044 volume of IEEE Std 1003.1-200x, Section 3.117, CPU Time (Execution Time). The mechanism  
2045 used to measure execution time is described in the Base Definitions volume of  
2046 IEEE Std 1003.1-200x, Section 4.9, Measurement of Execution Time.

2047 CPT If `_POSIX_CPUTIME` is defined, the following constant of the type `clockid_t` is defined in  
2048 `<time.h>`:

2049 `CLOCK_PROCESS_CPUTIME_ID`

2050 When this value of the type `clockid_t` is used in a `clock()` or `timer*()` function call, it is  
2051 interpreted as the identifier of the CPU-time clock associated with the process making the  
2052 function call.  
2053

2054 TCT If `_POSIX_THREAD_CPUTIME` is defined, the following constant of the type `clockid_t` is  
2055 defined in `<time.h>`:

2056 `CLOCK_THREAD_CPUTIME_ID`

2057 When this value of the type `clockid_t` is used in a `clock()` or `timer*()` function call, it is  
2058 interpreted as the identifier of the CPU-time clock associated with the thread making the  
2059 function call.

2060 **2.9 Threads**

2061 THR The functionality described in this section is dependent on support of the Threads option (and  
2062 the rest of this section is not further shaded for this option).

2063 This section defines functionality to support multiple flows of control, called *threads*, within a  
2064 process. For the definition of *threads*, see the Base Definitions volume of IEEE Std 1003.1-200x,  
2065 Section 3.393, Thread.

2066 The specific functional areas covered by threads and their scope include:

- 2067 • Thread management: the creation, control, and termination of multiple flows of control in the  
2068 same process under the assumption of a common shared address space
- 2069 • Synchronization primitives optimized for tightly coupled operation of multiple control flows  
2070 in a common, shared address space

2071 **2.9.1 Thread-Safety**

2072 All functions defined by this volume of IEEE Std 1003.1-200x shall be thread-safe, except that the  
2073 following functions<sup>1</sup> need not be thread-safe.

2074 \_\_\_\_\_

2075 1. The functions in the table are not shaded to denote applicable options. Individual reference pages should be consulted.

2076	<i>asctime()</i>	<i>ecvt()</i>	<i>gethostent()</i>	<i>getutxline()</i>	<i>putenv()</i>	
2077	<i>basename()</i>	<i>encrypt()</i>	<i>getlogin()</i>	<i>gmtime()</i>	<i>pututxline()</i>	
2078	<i>catgets()</i>	<i>endgrent()</i>	<i>getnetbyaddr()</i>	<i>hcreate()</i>	<i>rand()</i>	
2079	<i>crypt()</i>	<i>endpwent()</i>	<i>getnetbyname()</i>	<i>hdestroy()</i>	<i>readdir()</i>	
2080	<i>ctime()</i>	<i>endtxent()</i>	<i>getnetent()</i>	<i>hsearch()</i>	<i>setenv()</i>	
2081	<i>dbm_clearerr()</i>	<i>fcvt()</i>	<i>getopt()</i>	<i>inet_ntoa()</i>	<i>setgrent()</i>	
2082	<i>dbm_close()</i>	<i>ftw()</i>	<i>getprotobyname()</i>	<i>l64a()</i>	<i>setkey()</i>	
2083	<i>dbm_delete()</i>	<i>gcvt()</i>	<i>getprotobynumber()</i>	<i>lgamma()</i>	<i>setpwent()</i>	
2084	<i>dbm_error()</i>	<i>getc_unlocked()</i>	<i>getprotoent()</i>	<i>localeconv()</i>	<i>setutxent()</i>	
2085	<i>dbm_fetch()</i>	<i>getchar_unlocked()</i>	<i>getpwent()</i>	<i>localtime()</i>	<i>strerror()</i>	
2086	<i>dbm_firstkey()</i>	<i>getdate()</i>	<i>getpwnam()</i>	<i>lrand48()</i>	<i>strtok()</i>	
2087	<i>dbm_nextkey()</i>	<i>getenv()</i>	<i>getpwuid()</i>	<i>mrnd48()</i>	<i>ttyname()</i>	
2088	<i>dbm_open()</i>	<i>getgrent()</i>	<i>getservbyname()</i>	<i>nftw()</i>	<i>unsetenv()</i>	
2089	<i>dbm_store()</i>	<i>getgrgid()</i>	<i>getservbyport()</i>	<i>nl_langinfo()</i>	<i>wcstombs()</i>	
2090	<i>dirname()</i>	<i>getgrnam()</i>	<i>getservent()</i>	<i>ptsname()</i>	<i>wctomb()</i>	
2091	<i>derror()</i>	<i>gethostbyaddr()</i>	<i>getutxent()</i>	<i>putc_unlocked()</i>		
2092	<i>drand48()</i>	<i>gethostbyname()</i>	<i>getutxid()</i>	<i>putchar_unlocked()</i>		

2093 The *ctermid()* and *tmpnam()* functions need not be thread-safe if passed a NULL argument. The  
 2094 *wctomb()* and *wcrtombs()* functions need not be thread-safe if passed a NULL *ps* argument.

2095 Implementations shall provide internal synchronization as necessary in order to satisfy this  
 2096 requirement.

## 2097 2.9.2 Thread IDs

2098 Although implementations may have thread IDs that are unique in a system, applications  
 2099 should only assume that thread IDs are usable and unique within a single process. The effect of  
 2100 calling any of the functions defined in this volume of IEEE Std 1003.1-200x and passing as an  
 2101 argument the thread ID of a thread from another process is unspecified. A conforming  
 2102 implementation is free to reuse a thread ID after the thread terminates if it was created with the  
 2103 *detachstate* attribute set to *PTHREAD\_CREATE\_DETACHED* or if *pthread\_detach()* or  
 2104 *pthread\_join()* has been called for that thread. If a thread is detached, its thread ID is invalid for  
 2105 use as an argument in a call to *pthread\_detach()* or *pthread\_join()*.

## 2106 2.9.3 Thread Mutexes

2107 A thread that has blocked shall not prevent any unblocked thread that is eligible to use the same  
 2108 processing resources from eventually making forward progress in its execution. Eligibility for  
 2109 processing resources is determined by the scheduling policy.

2110 A thread shall become the owner of a mutex, *m*, when one of the following occurs: |

- 2111 • It returns successfully from *pthread\_mutex\_lock()* with *m* as the *mutex* argument.
- 2112 • It returns successfully from *pthread\_mutex\_trylock()* with *m* as the *mutex* argument.
- 2113 TMO • It returns successfully from *pthread\_mutex\_timedwait()* with *m* as the *mutex* argument.
- 2114 • It returns (successfully or not) from *pthread\_cond\_wait()* with *m* as the *mutex* argument  
 2115 (except as explicitly indicated otherwise for certain errors).
- 2116 • It returns (successfully or not) from *pthread\_cond\_timedwait()* with *m* as the *mutex* argument  
 2117 (except as explicitly indicated otherwise for certain errors).

2118 The thread shall remain the owner of *m* until one of the following occurs: |

- 2119           • It executes `pthread_mutex_unlock()` with *m* as the *mutex* argument
- 2120           • It blocks in a call to `pthread_cond_wait()` with *m* as the *mutex* argument.
- 2121           • It blocks in a call to `pthread_cond_timedwait()` with *m* as the *mutex* argument.
- 2122           The implementation shall behave as if at all times there is at most one owner of any mutex. |
- 2123           A thread that becomes the owner of a mutex is said to have *acquired* the mutex and the mutex is
- 2124           said to have become *locked*; when a thread gives up ownership of a mutex it is said to have
- 2125           *released* the mutex and the mutex is said to have become *unlocked*.
- 2126   **2.9.4 Thread Scheduling**
- 2127   TPS   The functionality described in this section is dependent on support of the Thread Execution
- 2128           Scheduling option (and the rest of this section is not further shaded for this option).
- 2129           **Thread Scheduling Attributes**
- 2130           In support of the scheduling function, threads have attributes which are accessed through the
- 2131           **pthread\_attr\_t** thread creation attributes object.
- 2132           The *contentionscope* attribute defines the scheduling contention scope of the thread to be either
- 2133           PTHREAD\_SCOPE\_PROCESS or PTHREAD\_SCOPE\_SYSTEM.
- 2134           The *inheritsched* attribute specifies whether a newly created thread is to inherit the scheduling
- 2135           attributes of the creating thread or to have its scheduling values set according to the other
- 2136           scheduling attributes in the **pthread\_attr\_t** object.
- 2137           The *schedpolicy* attribute defines the scheduling policy for the thread. The *schedparam* attribute
- 2138           defines the scheduling parameters for the thread. The interaction of threads having different
- 2139           policies within a process is described as part of the definition of those policies.
- 2140           If the Thread Execution Scheduling option is defined, and the *schedpolicy* attribute specifies one
- 2141           of the priority-based policies defined under this option, the *schedparam* attribute contains the
- 2142           scheduling priority of the thread. A conforming implementation ensures that the priority value
- 2143           in *schedparam* is in the range associated with the scheduling policy when the thread attributes
- 2144           object is used to create a thread, or when the scheduling attributes of a thread are dynamically
- 2145           modified. The meaning of the priority value in *schedparam* is the same as that of *priority*.
- 2146   TSP   If `_POSIX_THREAD_SPORADIC_SERVER` is defined, the *schedparam* attribute supports four
- 2147           new members that are used for the sporadic server scheduling policy. These members are
- 2148           *sched\_ss\_low\_priority*, *sched\_ss\_repl\_period*, *sched\_ss\_init\_budget*, and *sched\_ss\_max\_repl*. The
- 2149           meaning of these attributes is the same as in the definitions that appear under Section 2.8.4 (on
- 2150           page 494).
- 2151           When a process is created, its single thread has a scheduling policy and associated attributes
- 2152           equal to the process' policy and attributes. The default scheduling contention scope value is
- 2153           implementation-defined. The default values of other scheduling attributes are implementation-
- 2154           defined.

**2155 Thread Scheduling Contention Scope**

2156 The scheduling contention scope of a thread defines the set of threads with which the thread  
2157 competes for use of the processing resources. The scheduling operation selects at most one  
2158 thread to execute on each processor at any point in time and the thread's scheduling attributes  
2159 (for example, *priority*), whether under process scheduling contention scope or system scheduling  
2160 contention scope, are the parameters used to determine the scheduling decision.

2161 The scheduling contention scope, in the context of scheduling a mixed scope environment,  
2162 affects threads as follows:

- 2163 • A thread created with PTHREAD\_SCOPE\_SYSTEM scheduling contention scope contends  
2164 for resources with all other threads in the same scheduling allocation domain relative to their  
2165 system scheduling attributes. The system scheduling attributes of a thread created with  
2166 PTHREAD\_SCOPE\_SYSTEM scheduling contention scope are the scheduling attributes with  
2167 which the thread was created. The system scheduling attributes of a thread created with  
2168 PTHREAD\_SCOPE\_PROCESS scheduling contention scope are the implementation-defined  
2169 mapping into system attribute space of the scheduling attributes with which the thread was  
2170 created.
- 2171 • Threads created with PTHREAD\_SCOPE\_PROCESS scheduling contention scope contend  
2172 directly with other threads within their process that were created with  
2173 PTHREAD\_SCOPE\_PROCESS scheduling contention scope. The contention is resolved  
2174 based on the threads' scheduling attributes and policies. It is unspecified how such threads  
2175 are scheduled relative to threads in other processes or threads with  
2176 PTHREAD\_SCOPE\_SYSTEM scheduling contention scope.
- 2177 • Conforming implementations shall support the PTHREAD\_SCOPE\_PROCESS scheduling  
2178 contention scope, the PTHREAD\_SCOPE\_SYSTEM scheduling contention scope, or both.

**2179 Scheduling Allocation Domain**

2180 Implementations shall support scheduling allocation domains containing one or more  
2181 processors. It should be noted that the presence of multiple processors does not automatically  
2182 indicate a scheduling allocation domain size greater than one. Conforming implementations on  
2183 multi-processors may map all or any subset of the CPUs to one or multiple scheduling allocation  
2184 domains, and could define these scheduling allocation domains on a per-thread, per-process, or  
2185 per-system basis, depending on the types of applications intended to be supported by the  
2186 implementation. The scheduling allocation domain is independent of scheduling contention  
2187 scope, as the scheduling contention scope merely defines the set of threads with which a thread  
2188 contends for processor resources, while scheduling allocation domain defines the set of  
2189 processors for which it contends. The semantics of how this contention is resolved among  
2190 threads for processors is determined by the scheduling policies of the threads.

2191 The choice of scheduling allocation domain size and the level of application control over  
2192 scheduling allocation domains is implementation-defined. Conforming implementations may  
2193 change the size of scheduling allocation domains and the binding of threads to scheduling  
2194 allocation domains at any time.

2195 For application threads with scheduling allocation domains of size equal to one, the scheduling  
2196 rules defined for SCHED\_FIFO and SCHED\_RR shall be used; see **Scheduling Policies** (on page  
2197 494). All threads with system scheduling contention scope, regardless of the processes in which  
2198 they reside, compete for the processor according to their priorities. Threads with process  
2199 scheduling contention scope compete only with other threads with process scheduling  
2200 contention scope within their process.

2201 For application threads with scheduling allocation domains of size greater than one, the rules  
 2202 TSP defined for SCHED\_FIFO, SCHED\_RR, and SCHED\_SPORADIC shall be used in an  
 2203 implementation-defined manner. Each thread with system scheduling contention scope  
 2204 competes for the processors in its scheduling allocation domain in an implementation-defined  
 2205 manner according to its priority. Threads with process scheduling contention scope are  
 2206 scheduled relative to other threads within the same scheduling contention scope in the process.

2207 TSP If \_POSIX\_THREAD\_SPORADIC\_SERVER is defined, the rules defined for SCHED\_SPORADIC  
 2208 in **Scheduling Policies** (on page 494) shall be used in an implementation-defined manner for  
 2209 application threads whose scheduling allocation domain size is greater than one.

## 2210 Scheduling Documentation

2211 If \_POSIX\_PRIORITY\_SCHEDULING is defined, then any scheduling policies beyond  
 2212 TSP SCHED\_OTHER, SCHED\_FIFO, SCHED\_RR, and SCHED\_SPORADIC, as well as the effects of  
 2213 the scheduling policies indicated by these other values, and the attributes required in order to  
 2214 support such a policy, are implementation-defined. Furthermore, the implementation shall  
 2215 document the effect of all processor scheduling allocation domain values supported for these  
 2216 policies.

## 2217 2.9.5 Thread Cancellation

2218 The thread cancellation mechanism allows a thread to terminate the execution of any other  
 2219 thread in the process in a controlled manner. The target thread (that is, the one that is being  
 2220 canceled) is allowed to hold cancellation requests pending in a number of ways and to perform  
 2221 application-specific cleanup processing when the notice of cancellation is acted upon.

2222 Cancellation is controlled by the cancellation control functions. Each thread maintains its own  
 2223 cancelability state. Cancellation may only occur at cancellation points or when the thread is  
 2224 asynchronously cancelable.

2225 The thread cancellation mechanism described in this section depends upon programs having set  
 2226 *deferred cancelability* state, which is specified as the default. Applications shall also carefully  
 2227 follow static lexical scoping rules in their execution behavior. For example, use of *setjmp()*,  
 2228 *return*, *goto*, and so on, to leave user-defined cancellation scopes without doing the necessary  
 2229 scope pop operation results in undefined behavior.

2230 Use of asynchronous cancelability while holding resources which potentially need to be released  
 2231 may result in resource loss. Similarly, cancellation scopes may only be safely manipulated  
 2232 (pushed and popped) when the thread is in the *deferred* or *disabled* cancelability states.

### 2233 2.9.5.1 Cancelability States

2234 The cancelability state of a thread determines the action taken upon receipt of a cancellation  
 2235 request. The thread may control cancellation in a number of ways.

2236 Each thread maintains its own cancelability state, which may be encoded in two bits:

2237 1. Cancelability-Enable: When cancelability is PTHREAD\_CANCEL\_DISABLE (as defined in  
 2238 the Base Definitions volume of IEEE Std 1003.1-200x, <**pthread.h**>), cancellation requests  
 2239 against the target thread are held pending. By default, cancelability is set to  
 2240 PTHREAD\_CANCEL\_ENABLE (as defined in <**pthread.h**>).

2241 2. Cancelability Type: When cancelability is enabled and the cancelability type is  
 2242 PTHREAD\_CANCEL\_ASYNCHRONOUS (as defined in <**pthread.h**>), new or pending  
 2243 cancellation requests may be acted upon at any time. When cancelability is enabled and the  
 2244 cancelability type is PTHREAD\_CANCEL\_DEFERRED (as defined in <**pthread.h**>),

2245 cancellation requests are held pending until a cancellation point (see below) is reached. If  
 2246 cancelability is disabled, the setting of the cancelability type has no immediate effect as all  
 2247 cancellation requests are held pending; however, once cancelability is enabled again the  
 2248 new type is in effect. The cancelability type is PTHREAD\_CANCEL\_DEFERRED in all  
 2249 newly created threads including the thread in which *main()* was first invoked.

### 2250 2.9.5.2 Cancellation Points

2251 Cancellation points shall occur when a thread is executing the following functions: |

2252	<i>accept()</i>	<i>mq_timedsend()</i>	<i>putpmsg()</i>	<i>sigsuspend()</i>
2253	<i>aio_suspend()</i>	<i>msgrcv()</i>	<i>pwrite()</i>	<i>sigtimedwait()</i>
2254	<i>clock_nanosleep()</i>	<i>msgsnd()</i>	<i>read()</i>	<i>sigwait()</i>
2255	<i>close()</i>	<i>msync()</i>	<i>readv()</i>	<i>sigwaitinfo()</i>
2256	<i>connect()</i>	<i>nanosleep()</i>	<i>recv()</i>	<i>sleep()</i>
2257	<i>creat()</i>	<i>open()</i>	<i>recvfrom()</i>	<i>system()</i>
2258	<i>fcntl()</i> <sup>2</sup>	<i>pause()</i>	<i>recvmsg()</i>	<i>tcdrain()</i>
2259	<i>fsync()</i>	<i>poll()</i>	<i>select()</i>	<i>usleep()</i>
2260	<i>getmsg()</i>	<i>pread()</i>	<i>sem_timedwait()</i>	<i>wait()</i>
2261	<i>getpmsg()</i>	<i>pthread_cond_timedwait()</i>	<i>sem_wait()</i>	<i>waitid()</i>
2262	<i>lockf()</i>	<i>pthread_cond_wait()</i>	<i>send()</i>	<i>waitpid()</i>
2263	<i>mq_receive()</i>	<i>pthread_join()</i>	<i>sendmsg()</i>	<i>write()</i>
2264	<i>mq_send()</i>	<i>pthread_testcancel()</i>	<i>sendto()</i>	<i>writev()</i>
2265	<i>mq_timedreceive()</i>	<i>putmsg()</i>	<i>sigpause()</i>	

2266 \_\_\_\_\_

2267 2. When the *cmd* argument is F\_SETLKW.

2268 A cancelation point may also occur when a thread is executing the following functions:

2269	<i>catclose()</i>	<i>ftell()</i>	<i>getwc()</i>	<i>pthread_rwlock_wrlock()</i>
2270	<i>catgets()</i>	<i>ftello()</i>	<i>getwchar()</i>	<i>putc()</i>
2271	<i>catopen()</i>	<i>ftw()</i>	<i>getwd()</i>	<i>putc_unlocked()</i>
2272	<i>closedir()</i>	<i>fwprintf()</i>	<i>glob()</i>	<i>putchar()</i>
2273	<i>closelog()</i>	<i>fwrite()</i>	<i>iconv_close()</i>	<i>putchar_unlocked()</i>
2274	<i>ctermid()</i>	<i>fwscanf()</i>	<i>iconv_open()</i>	<i>puts()</i>
2275	<i>dbm_close()</i>	<i>getc()</i>	<i>ioctl()</i>	<i>pututxline()</i>
2276	<i>dbm_delete()</i>	<i>getc_unlocked()</i>	<i>lseek()</i>	<i>putwc()</i>
2277	<i>dbm_fetch()</i>	<i>getchar()</i>	<i>mkstemp()</i>	<i>putwchar()</i>
2278	<i>dbm_nextkey()</i>	<i>getchar_unlocked()</i>	<i>nftw()</i>	<i>readdir()</i>
2279	<i>dbm_open()</i>	<i>getcwd()</i>	<i>opendir()</i>	<i>readdir_r()</i>
2280	<i>dbm_store()</i>	<i>getdate()</i>	<i>openlog()</i>	<i>remove()</i>
2281	<i>dlclose()</i>	<i>getgrent()</i>	<i>pclose()</i>	<i>rename()</i>
2282	<i>dlopen()</i>	<i>getgrgid()</i>	<i>perror()</i>	<i>rewind()</i>
2283	<i>endgrent()</i>	<i>getgrgid_r()</i>	<i>popen()</i>	<i>rewinddir()</i>
2284	<i>endhostent()</i>	<i>getgrnam()</i>	<i>posix_fadvise()</i>	<i>scanf()</i>
2285	<i>endnetent()</i>	<i>getgrnam_r()</i>	<i>posix_fallocate()</i>	<i>seekdir()</i>
2286	<i>endprotoent()</i>	<i>gethostbyaddr()</i>	<i>posix_madvise()</i>	<i>semop()</i>
2287	<i>endpwent()</i>	<i>gethostbyname()</i>	<i>posix_spawn()</i>	<i>setgrent()</i>
2288	<i>endservent()</i>	<i>gethostent()</i>	<i>posix_spawnnp()</i>	<i>sethostent()</i>
2289	<i>endutxent()</i>	<i>gethostname()</i>	<i>posix_trace_clear()</i>	<i>setnetent()</i>
2290	<i>fclose()</i>	<i>getlogin()</i>	<i>posix_trace_close()</i>	<i>setprotoent()</i>
2291	<i>fcntl()</i> <sup>3</sup>	<i>getlogin_r()</i>	<i>posix_trace_create()</i>	<i>setpwent()</i>
2292	<i>fflush()</i>	<i>getnetbyaddr()</i>	<i>posix_trace_create_withlog()</i>	<i>setservent()</i>
2293	<i>fgetc()</i>	<i>getnetbyname()</i>	<i>posix_trace_eventtypelist_getnext_id()</i>	<i>setutxent()</i>
2294	<i>fgetpos()</i>	<i>getnetent()</i>	<i>posix_trace_eventtypelist_rewind()</i>	<i>strerror()</i>
2295	<i>fgets()</i>	<i>getprotobyname()</i>	<i>posix_trace_flush()</i>	<i>syslog()</i>
2296	<i>fgetwc()</i>	<i>getprotobynumber()</i>	<i>posix_trace_get_attr()</i>	<i>tmpfile()</i>
2297	<i>fgetws()</i>	<i>getprotoent()</i>	<i>posix_trace_get_filter()</i>	<i>tmpnam()</i>
2298	<i>fopen()</i>	<i>getpwent()</i>	<i>posix_trace_get_status()</i>	<i>ttyname()</i>
2299	<i>fprintf()</i>	<i>getpwnam()</i>	<i>posix_trace_getnext_event()</i>	<i>ttyname_r()</i>
2300	<i>fputc()</i>	<i>getpwnam_r()</i>	<i>posix_trace_open()</i>	<i>ungetc()</i>
2301	<i>fputs()</i>	<i>getpwuid()</i>	<i>posix_trace_rewind()</i>	<i>ungetwc()</i>
2302	<i>fputwc()</i>	<i>getpwuid_r()</i>	<i>posix_trace_set_filter()</i>	<i>unlink()</i>
2303	<i>fputws()</i>	<i>gets()</i>	<i>posix_trace_shutdown()</i>	<i>vfprintf()</i>
2304	<i>fread()</i>	<i>getservbyname()</i>	<i>posix_trace_timedgetnext_event()</i>	<i>vwprintf()</i>
2305	<i>freopen()</i>	<i>getservbyport()</i>	<i>posix_typed_mem_open()</i>	<i>vprintf()</i>
2306	<i>fscanf()</i>	<i>getservent()</i>	<i>printf()</i>	<i>vwprintf()</i>
2307	<i>fseek()</i>	<i>getutxent()</i>	<i>pthread_rwlock_rdlock()</i>	<i>wprintf()</i>
2308	<i>fseeko()</i>	<i>getutxid()</i>	<i>pthread_rwlock_timedrdlock()</i>	<i>wscanf()</i>
2309	<i>fsetpos()</i>	<i>getutxline()</i>	<i>pthread_rwlock_timedwrlock()</i>	

2310 An implementation shall not introduce cancelation points into any other functions specified in  
2311 this volume of IEEE Std 1003.1-200x.

2312 \_\_\_\_\_

2313 3. For any value of the *cmd* argument.



2314 The side effects of acting upon a cancelation request while suspended during a call of a function  
2315 are the same as the side effects that may be seen in a single-threaded program when a call to a  
2316 function is interrupted by a signal and the given function returns [EINTR]. Any such side effects  
2317 occur before any cancelation cleanup handlers are called.

2318 Whenever a thread has cancelability enabled and a cancelation request has been made with that  
2319 thread as the target, and the thread then calls any function that is a cancelation point (such as  
2320 *pthread\_testcancel()* or *read()*), the cancelation request shall be acted upon before the function  
2321 returns. If a thread has cancelability enabled and a cancelation request is made with the thread  
2322 as a target while the thread is suspended at a cancelation point, the thread shall be awakened  
2323 and the cancelation request shall be acted upon. However, if the thread is suspended at a  
2324 cancelation point and the event for which it is waiting occurs before the cancelation request is  
2325 acted upon, it is unspecified whether the cancelation request is acted upon or whether the  
2326 cancelation request remains pending and the thread resumes normal execution.

### 2327 2.9.5.3 Thread Cancelation Cleanup Handlers

2328 Each thread maintains a list of cancelation cleanup handlers. The programmer uses the  
2329 *pthread\_cleanup\_push()* and *pthread\_cleanup\_pop()* functions to place routines on and remove  
2330 routines from this list.

2331 When a cancelation request is acted upon, the routines in the list are invoked one by one in LIFO  
2332 sequence; that is, the last routine pushed onto the list (Last In) is the first to be invoked (First  
2333 Out). The thread invokes the cancelation cleanup handler with cancelation disabled until the last  
2334 cancelation cleanup handler returns. When the cancelation cleanup handler for a scope is  
2335 invoked, the storage for that scope remains valid. If the last cancelation cleanup handler returns,  
2336 thread execution is terminated and a status of `PTHREAD_CANCELED` is made available to any  
2337 threads joining with the target. The symbolic constant `PTHREAD_CANCELED` expands to a  
2338 constant expression of type (**void \***) whose value matches no pointer to an object in memory nor  
2339 the value `NULL`.

2340 The cancelation cleanup handlers are also invoked when the thread calls *pthread\_exit()*.

2341 A side effect of acting upon a cancelation request while in a condition variable wait is that the  
2342 mutex is re-acquired before calling the first cancelation cleanup handler. In addition, the thread  
2343 is no longer considered to be waiting for the condition and the thread shall not have consumed  
2344 any pending condition signals on the condition.

2345 A cancelation cleanup handler cannot exit via *longjmp()* or *siglongjmp()*.

### 2346 2.9.5.4 Async-Cancel Safety

2347 The *pthread\_cancel()*, *pthread\_setcancelstate()*, and *pthread\_setcanceltype()* functions are defined to  
2348 be async-cancel safe.

2349 No other functions in this volume of IEEE Std 1003.1-200x are required to be async-cancel-safe.

**2350 2.9.6 Thread Read-Write Locks**

2351 Multiple readers, single writer (read-write) locks allow many threads to have simultaneous  
2352 read-only access to data while allowing only one thread to have exclusive write access at any  
2353 given time. They are typically used to protect data that is read more frequently than it is  
2354 changed.

2355 One or more readers acquire read access to the resource by performing a read lock operation on  
2356 the associated read-write lock. A writer acquires exclusive write access by performing a write  
2357 lock operation. Basically, all readers exclude any writers and a writer excludes all readers and  
2358 any other writers.

2359 A thread that has blocked on a read-write lock (for example, has not yet returned from a  
2360 *pthread\_rwlock\_rdlock()* or *pthread\_rwlock\_wrlock()* call) shall not prevent any unblocked thread  
2361 that is eligible to use the same processing resources from eventually making forward progress in  
2362 its execution. Eligibility for processing resources shall be determined by the scheduling policy.

2363 Read-write locks can be used to synchronize threads in the current process and other processes if  
2364 they are allocated in memory that is writable and shared among the cooperating processes and  
2365 have been initialized for this behavior.

**2366 2.9.7 Thread Interactions with Regular File Operations**

2367 All of the functions *chmod()*, *close()*, *fchmod()*, *fcntl()*, *fstat()*, *ftruncate()*, *lseek()*, *open()*, *read()*,  
2368 *readlink()*, *stat()*, *symlink()*, and *write()* shall be atomic with respect to each other in the effects  
2369 specified in IEEE Std 1003.1-200x when they operate on regular files. If two threads each call one  
2370 of these functions, each call shall either see all of the specified effects of the other call, or none of  
2371 them.

**2372 2.10 Sockets**

2373 A socket is an endpoint for communication using the facilities described in this section. A socket  
2374 is created with a specific socket type, described in Section 2.10.6 (on page 509), and is associated  
2375 with a specific protocol, detailed in Section 2.10.3 (on page 509). A socket is accessed via a file  
2376 descriptor obtained when the socket is created. |

**2377 2.10.1 Address Families** |

2378 All network protocols are associated with a specific address family. An address family provides |  
2379 basic services to the protocol implementation to allow it to function within a specific network |  
2380 environment. These services may include packet fragmentation and reassembly, routing, |  
2381 addressing, and basic transport. An address family is normally comprised of a number of |  
2382 protocols, one per socket type. Each protocol is characterized by an abstract socket type. It is not |  
2383 required that an address family support all socket types. An address family may contain |  
2384 multiple protocols supporting the same socket abstraction. |

2385 Section 2.10.17 (on page 516), Section 2.10.19 (on page 517), and Section 2.10.20 (on page 517),  
2386 respectively, describe the use of sockets for local UNIX connections, for Internet protocols based  
2387 on IPv4, and for Internet protocols based on IPv6. |

2388 **2.10.2 Addressing**

2389 An address family defines the format of a socket address. All network addresses are described  
 2390 using a general structure, called a **sockaddr**, as defined in the Base Definitions volume of  
 2391 IEEE Std 1003.1-200x, `<sys/socket.h>`. However, each address family imposes finer and more  
 2392 specific structure, generally defining a structure with fields specific to the address family. The  
 2393 field *sa\_family* in the **sockaddr** structure contains the address family identifier, specifying the  
 2394 format of the *sa\_data* area. The size of the *sa\_data* area is unspecified.

2395 **2.10.3 Protocols**

2396 A protocol supports one of the socket abstractions detailed in Section 2.10.6. Selecting a protocol  
 2397 involves specifying the address family, socket type, and protocol number to the *socket()*  
 2398 function. Certain semantics of the basic socket abstractions are protocol-specific. All protocols  
 2399 are expected to support the basic model for their particular socket type, but may, in addition,  
 2400 provide non-standard facilities or extensions to a mechanism.

2401 **2.10.4 Routing**

2402 Sockets provides packet routing facilities. A routing information database is maintained, which  
 2403 is used in selecting the appropriate network interface when transmitting packets.

2404 **2.10.5 Interfaces**

2405 Each network interface in a system corresponds to a path through which messages can be sent  
 2406 and received. A network interface usually has a hardware device associated with it, though  
 2407 certain interfaces such as the loopback interface, do not.

2408 **2.10.6 Socket Types**

2409 A socket is created with a specific type, which defines the communication semantics and which  
 2410 RS allows the selection of an appropriate communication protocol. Four types are defined:  
 2411 `SOCK_RAW`, `SOCK_STREAM`, `SOCK_SEQPACKET`, and `SOCK_DGRAM`. Implementations  
 2412 may specify additional socket types.

2413 The `SOCK_STREAM` socket type provides reliable, sequenced, full-duplex octet streams  
 2414 between the socket and a peer to which the socket is connected. A socket of type  
 2415 `SOCK_STREAM` must be in a connected state before any data may be sent or received. Record  
 2416 boundaries are not maintained; data sent on a stream socket using output operations of one size  
 2417 may be received using input operations of smaller or larger sizes without loss of data. Data may  
 2418 be buffered; successful return from an output function does not imply that the data has been  
 2419 delivered to the peer or even transmitted from the local system. If data cannot be successfully  
 2420 transmitted within a given time then the connection is considered broken, and subsequent  
 2421 operations shall fail. A `SIGPIPE` signal is raised if a thread sends on a broken stream (one that is  
 2422 no longer connected). Support for an out-of-band data transmission facility is protocol-specific.

2423 The `SOCK_SEQPACKET` socket type is similar to the `SOCK_STREAM` type, and is also  
 2424 connection-oriented. The only difference between these types is that record boundaries are  
 2425 maintained using the `SOCK_SEQPACKET` type. A record can be sent using one or more output  
 2426 operations and received using one or more input operations, but a single operation never  
 2427 transfers parts of more than one record. Record boundaries are visible to the receiver via the  
 2428 `MSG_EOR` flag in the received message flags returned by the *recvmsg()* function. It is protocol-  
 2429 specific whether a maximum record size is imposed.

2430 The `SOCK_DGRAM` socket type supports connectionless data transfer which is not necessarily  
 2431 acknowledged or reliable. Datagrams may be sent to the address specified (possibly multicast or

2432 broadcast) in each output operation, and incoming datagrams may be received from multiple |  
2433 sources. The source address of each datagram is available when receiving the datagram. An |  
2434 application may also pre-specify a peer address, in which case calls to output functions shall |  
2435 send to the pre-specified peer. If a peer has been specified, only datagrams from that peer shall |  
2436 be received. A datagram must be sent in a single output operation, and must be received in a |  
2437 single input operation. The maximum size of a datagram is protocol-specific; with some |  
2438 protocols, the limit is implementation-defined. Output datagrams may be buffered within the |  
2439 system; thus, a successful return from an output function does not guarantee that a datagram is |  
2440 actually sent or received. However, implementations should attempt to detect any errors |  
2441 possible before the return of an output function, reporting any error by an unsuccessful return |  
2442 value.

2443 RS The SOCK\_RAW socket type is similar to the SOCK\_DGRAM type. It differs in that it is |  
2444 normally used with communication providers that underlie those used for the other socket |  
2445 types. For this reason, the creation of a socket with type SOCK\_RAW shall require appropriate |  
2446 privilege. The format of datagrams sent and received with this socket type generally include |  
2447 specific protocol headers, and the formats are protocol-specific and implementation-defined.

### 2448 **2.10.7 Socket I/O Mode**

2449 The I/O mode of a socket is described by the O\_NONBLOCK file status flag which pertains to |  
2450 the open file description for the socket. This flag is initially off when a socket is created, but may |  
2451 be set and cleared by the use of the F\_SETFL command of the *fcntl()* function.

2452 When the O\_NONBLOCK flag is set, functions that would normally block until they are |  
2453 complete shall either return immediately with an error, or shall complete asynchronously to the |  
2454 execution of the calling process. Data transfer operations (the *read()*, *write()*, *send()*, and *recv()* |  
2455 functions) shall complete immediately, transfer only as much as is available, and then return |  
2456 without blocking, or return an error indicating that no transfer could be made without blocking. |  
2457 The *connect()* function initiates a connection and shall return without blocking when |  
2458 O\_NONBLOCK is set; it shall return the error [EINPROGRESS] to indicate that the connection |  
2459 was initiated successfully, but that it has not yet completed.

### 2460 **2.10.8 Socket Owner**

2461 The owner of a socket is unset when a socket is created. The owner may be set to a process ID or |  
2462 process group ID using the F\_SETOWN command of the *fcntl()* function.

### 2463 **2.10.9 Socket Queue Limits**

2464 The transmit and receive queue sizes for a socket are set when the socket is created. The default |  
2465 sizes used are both protocol-specific and implementation-defined. The sizes may be changed |  
2466 using the *setsockopt()* function.

### 2467 **2.10.10 Pending Error**

2468 Errors may occur asynchronously, and be reported to the socket in response to input from the |  
2469 network protocol. The socket stores the pending error to be reported to a user of the socket at the |  
2470 next opportunity. The error is returned in response to a subsequent *send()*, *recv()*, or *getsockopt()* |  
2471 operation on the socket, and the pending error is then cleared.

**2.10.11 Socket Receive Queue**

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

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

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

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

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

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

**2517 2.10.13 Connection Indication Queue**

2518 Sockets that are used to accept incoming connections maintain a queue of outstanding  
2519 connection indications. This queue is a list of connections that are awaiting acceptance by the  
2520 application; see *listen()*.

**2521 2.10.14 Signals**

2522 One category of event at the socket interface is the generation of signals. These signals report  
2523 protocol events or process errors relating to the state of the socket. The generation or delivery of  
2524 a signal does not change the state of the socket, although the generation of the signal may have  
2525 been caused by a state change.

2526 The SIGPIPE signal shall be sent to a thread that attempts to send data on a socket that is no  
2527 longer able to send. In addition, the send operation fails with the error [EPIPE].

2528 If a socket has an owner, the SIGURG signal is sent to the owner of the socket when it is notified  
2529 of expedited or out-of-band data. The socket state at this time is protocol-dependent, and the  
2530 status of the socket is specified in Section 2.10.17 (on page 516), Section 2.10.19 (on page 517),  
2531 and Section 2.10.20 (on page 517). Depending on the protocol, the expedited data may or may  
2532 not have arrived at the time of signal generation.

**2533 2.10.15 Asynchronous Errors**

2534 If any of the following conditions occur asynchronously for a socket, the corresponding value  
2535 listed below shall become the pending error for the socket:

2536 [ECONNABORTED]

2537 The connection was aborted locally.

2538 [ECONNREFUSED]

2539 For a connection-mode socket attempting a non-blocking connection, the attempt to connect  
2540 was forcefully rejected. For a connectionless-mode socket, an attempt to deliver a datagram  
2541 was forcefully rejected.

2542 [ECONNRESET]

2543 The peer has aborted the connection.

2544 [EHOSTDOWN]

2545 The destination host has been determined to be down or disconnected.

2546 [EHOSTUNREACH]

2547 The destination host is not reachable.

2548 [EMSGSIZE]

2549 For a connectionless-mode socket, the size of a previously sent datagram prevented  
2550 delivery.

2551 [ENETDOWN]

2552 The local network connection is not operational.

2553 [ENETRESET]

2554 The connection was aborted by the network.

2555 [ENETUNREACH]

2556 The destination network is not reachable.

2557 **2.10.16 Use of Options**

2558 There are a number of socket options which either specialize the behavior of a socket or provide  
 2559 useful information. These options may be set at different protocol levels and are always present  
 2560 at the uppermost “socket” level.

2561 Socket options are manipulated by two functions, *getsockopt()* and *setsockopt()*. These functions  
 2562 allow an application program to customize the behavior and characteristics of a socket to  
 2563 provide the desired effect.

2564 All of the options have default values. The type and meaning of these values is defined by the  
 2565 protocol level to which they apply. Instead of using the default values, an application program  
 2566 may choose to customize one or more of the options. However, in the bulk of cases, the default  
 2567 values are sufficient for the application.

2568 Some of the options are used to enable or disable certain behavior within the protocol modules  
 2569 (for example, turn on debugging) while others may be used to set protocol-specific information  
 2570 (for example, IP time-to-live on all the application’s outgoing packets). As each of the options is  
 2571 introduced, its effect on the underlying protocol modules is described.

2572 Table 2-1 shows the value for the socket level.

2573 **Table 2-1** Value of Level for Socket Options

Name	Description
SOL_SOCKET	Options are intended for the sockets level.

2576 Table 2-2 (on page 514) lists those options present at the socket level; that is, when the *level*  
 2577 parameter of the *getsockopt()* or *setsockopt()* function is SOL\_SOCKET, the types of the option  
 2578 value parameters associated with each option, and a brief synopsis of the meaning of the option  
 2579 value parameter. Unless otherwise noted, each may be examined with *getsockopt()* and set with  
 2580 *setsockopt()* on all types of socket.

Table 2-2 Socket-Level Options

Option	Parameter Type	Parameter Meaning
SO_BROADCAST	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).

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.

The SO\_DEBUG option 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.

The SO\_DONTROUTE option 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.

The SO\_ERROR option is used only on *getsockopt()*. When this option is specified, *getsockopt()* shall return any pending error on the socket and clear the error status. It shall return 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.



2628 The SO\_KEEPALIVE option enables the periodic transmission of messages on a connected |  
2629 socket. The behavior of this option is protocol-specific. The default value for SO\_KEEPALIVE is |  
2630 zero, specifying that this capability is turned off. |

2631 The SO\_LINGER option controls the action of the interface when unsent messages are queued |  
2632 on a socket and a *close()* is performed. The details of this option are protocol-specific. The |  
2633 default value for SO\_LINGER is zero, or off, for the *L\_onoff* element of the option value and zero |  
2634 seconds for the linger time specified by the *L\_linger* element.

2635 The SO\_OOBINLINE option is valid only on protocols that support out-of-band data. The |  
2636 SO\_OOBINLINE option requests that out-of-band data be placed in the normal data input queue |  
2637 as received; it is then accessible using the *read()* or *recv()* functions without the MSG\_OOB flag |  
2638 set. The default for SO\_OOBINLINE is off; that is, for out-of-band data not to be placed in the |  
2639 normal data input queue.

2640 The SO\_RCVBUF option requests that the buffer space allocated for receive operations on this |  
2641 socket be set to the value, in bytes, of the option value. Applications may wish to increase buffer |  
2642 size for high volume connections, or may decrease buffer size to limit the possible backlog of |  
2643 incoming data. The default value for the SO\_RCVBUF option value is implementation-defined, |  
2644 and may vary by protocol. |

2645 The maximum value for the option for a socket may be obtained by the use of the *fpathconf()* |  
2646 function, using the value *\_PC\_SOCKET\_MAXBUF*.

2647 The SO\_RCVLOWAT option sets the minimum number of bytes to process for socket input |  
2648 operations. In general, receive calls block until any (non-zero) amount of data is received, then |  
2649 return the smaller of the amount available or the amount requested. The default value for |  
2650 SO\_RCVLOWAT is 1, and does not affect the general case. If SO\_RCVLOWAT is set to a larger |  
2651 value, blocking receive calls normally wait until they have received the smaller of the low water |  
2652 mark value or the requested amount. Receive calls may still return less than the low water mark |  
2653 if an error occurs, a signal is caught, or the type of data next in the receive queue is different |  
2654 from that returned (for example, out-of-band data). As mentioned previously, the default value |  
2655 for SO\_RCVLOWAT is 1 byte. It is implementation-defined whether the SO\_RCVLOWAT option |  
2656 can be set. |

2657 The SO\_RCVTIMEO option is an option to set a timeout value for input operations. It accepts a |  
2658 **timeval** structure with the number of seconds and microseconds specifying the limit on how |  
2659 long to wait for an input operation to complete. If a receive operation has blocked for this much |  
2660 time without receiving additional data, it shall return with a partial count or *errno* shall be set to |  
2661 [EWOULDBLOCK] if no data were received. The default for this option is the value zero, which |  
2662 indicates that a receive operation will not timeout. It is implementation-defined whether the |  
2663 SO\_RCVTIMEO option can be set.

2664 The SO\_REUSEADDR option indicates that the rules used in validating addresses supplied in a |  
2665 *bind()* should allow reuse of local addresses. Operation of this option is protocol-specific. The |  
2666 default value for SO\_REUSEADDR is off; that is, reuse of local addresses is not permitted.

2667 The SO\_SNDBUF option requests that the buffer space allocated for send operations on this |  
2668 socket be set to the value, in bytes, of the option value. The default value for the SO\_SNDBUF |  
2669 option value is implementation-defined, and may vary by protocol. The maximum value for the |  
2670 option for a socket may be obtained by the use of the *fpathconf()* function, using the value |  
2671 *\_PC\_SOCKET\_MAXBUF*.

2672 The SO\_SNDLOWAT option sets the minimum number of bytes to process for socket output |  
2673 operations. Most output operations process all of the data supplied by the call, delivering data to |  
2674 the protocol for transmission and blocking as necessary for flow control. Non-blocking output |  
2675 operations process as much data as permitted subject to flow control without blocking, but |

2676 process no data if flow control does not allow the smaller of the send low water mark value or  
 2677 the entire request to be processed. A *select()* operation testing the ability to write to a socket shall  
 2678 return true only if the send low water mark could be processed. The default value for  
 2679 SO\_SNDLOWAT is implementation-defined and protocol-specific. It is implementation-defined  
 2680 whether the SO\_SNDLOWAT option can be set.

2681 The SO\_SNDTIMEO option is an option to set a timeout value for the amount of time that an  
 2682 output function shall block because flow control prevents data from being sent. As noted in  
 2683 Table 2-2 (on page 514), the option value is a **timeval** structure with the number of seconds and  
 2684 microseconds specifying the limit on how long to wait for an output operation to complete. If a  
 2685 send operation has blocked for this much time, it shall return with a partial count or *errno* set to  
 2686 [EWOULDBLOCK] if no data were sent. The default for this option is the value zero, which  
 2687 indicates that a send operation will not timeout. It is implementation-defined whether the  
 2688 SO\_SNDTIMEO option can be set.

2689 The SO\_TYPE option is used only on *getsockopt()*. When this option is specified, *getsockopt()*  
 2690 shall return the type of the socket (for example, SOCK\_STREAM). This option is useful to  
 2691 servers that inherit sockets on start-up. SO\_TYPE has no default value.

## 2692 2.10.17 Use of Sockets for Local UNIX Connections

2693 Support for UNIX domain sockets is mandatory.

2694 UNIX domain sockets provide process-to-process communication in a single system.

### 2695 2.10.17.1 Headers

2696 The symbolic constant AF\_UNIX defined in the `<sys/socket.h>` header is used to identify the  
 2697 UNIX domain address family. The `<sys/un.h>` header contains other definitions used in  
 2698 connection with UNIX domain sockets. See the Base Definitions volume of IEEE Std 1003.1-200x,  
 2699 Chapter 13, Headers.

2700 The **sockaddr\_storage** structure defined in `<sys/socket.h>` shall be large enough to  
 2701 accommodate a **sockaddr\_un** structure (see the `<sys/un.h>` header defined in the Base  
 2702 Definitions volume of IEEE Std 1003.1-200x, Chapter 13, Headers) and shall be aligned at an  
 2703 appropriate boundary so that pointers to it can be cast as pointers to **sockaddr\_un** structures  
 2704 and used to access the fields of those structures without alignment problems. When a  
 2705 **sockaddr\_storage** structure is cast as a **sockaddr\_un** structure, the *ss\_family* field maps onto the  
 2706 *sun\_family* field.

## 2707 2.10.18 Use of Sockets over Internet Protocols

2708 When a socket is created in the Internet family with a protocol value of zero, the implementation  
 2709 shall use the protocol listed below for the type of socket created.

2710 SOCK\_STREAM      IPPROTO\_TCP.

2711 SOCK\_DGRAM      IPPROTO\_UDP.

2712 RS      SOCK\_RAW      IPPROTO\_RAW.

2713 SOCK\_SEQPACKET    Unspecified.

2714 RS      A raw interface to IP is available by creating an Internet socket of type SOCK\_RAW. The default  
 2715 protocol for type SOCK\_RAW shall be identified in the IP header with the value  
 2716 IPPROTO\_RAW. Applications should not use the default protocol when creating a socket with  
 2717 type SOCK\_RAW, but should identify a specific protocol by value. The ICMP control protocol is  
 2718 accessible from a raw socket by specifying a value of IPPROTO\_ICMP for protocol.

## 2719 **2.10.19 Use of Sockets over Internet Protocols Based on IPv4**

2720 Support for sockets over Internet protocols based on IPv4 is mandatory.

### 2721 *2.10.19.1 Headers*

2722 The symbolic constant `AF_INET` defined in the `<sys/socket.h>` header is used to identify the  
 2723 IPv4 Internet address family. The `<netinet/in.h>` header contains other definitions used in  
 2724 connection with IPv4 Internet sockets. See the Base Definitions volume of IEEE Std 1003.1-200x,  
 2725 Chapter 13, Headers.

2726 The `sockaddr_storage` structure defined in `<sys/socket.h>` shall be large enough to  
 2727 accommodate a `sockaddr_in` structure (see the `<netinet/in.h>` header defined in the Base  
 2728 Definitions volume of IEEE Std 1003.1-200x, Chapter 13, Headers) and shall be aligned at an  
 2729 appropriate boundary so that pointers to it can be cast as pointers to `sockaddr_in` structures and  
 2730 used to access the fields of those structures without alignment problems. When a  
 2731 `sockaddr_storage` structure is cast as a `sockaddr_in` structure, the `ss_family` field maps onto the  
 2732 `sin_family` field.

## 2733 **2.10.20 Use of Sockets over Internet Protocols Based on IPv6**

2734 IP6 This section describes extensions to support sockets over Internet protocols based on IPv6. This  
 2735 functionality is dependent on support of the IPV6 option (and the rest of this section is not  
 2736 further shaded for this option).

2737 To enable smooth transition from IPv4 to IPv6, the features defined in this section may, in certain  
 2738 circumstances, also be used in connection with IPv4; see Section 2.10.20.2 (on page 518).

### 2739 *2.10.20.1 Addressing*

2740 IPv6 overcomes the addressing limitations of previous versions by using 128-bit addresses  
 2741 instead of 32-bit addresses. The IPv6 address architecture is described in RFC 2373.

2742 There are three kinds of IPv6 address:

#### 2743 Unicast

2744 Identifies a single interface.

2745 A unicast address can be global, link-local (designed for use on a single link), or site-local  
 2746 (designed for systems not connected to the Internet). Link-local and site-local addresses  
 2747 need not be globally unique.

#### 2748 Anycast

2749 Identifies a set of interfaces such that a packet sent to the address can be delivered to any  
 2750 member of the set.

2751 An anycast address is similar to a unicast address; the nodes to which an anycast address is  
 2752 assigned must be explicitly configured to know that it is an anycast address.

#### 2753 Multicast

2754 Identifies a set of interfaces such that a packet sent to the address should be delivered to  
 2755 every member of the set.

2756 An application can send multicast datagrams by simply specifying an IPv6 multicast  
 2757 address in the `address` argument of `sendto()`. To receive multicast datagrams, an application  
 2758 must join the multicast group (using `setsockopt()` with `IPV6_JOIN_GROUP`) and must bind  
 2759 to the socket the UDP port on which datagrams will be received. Some applications should  
 2760 also bind the multicast group address to the socket, to prevent other datagrams destined to  
 2761 that port from being delivered to the socket.

2762 A multicast address can be global, node-local, link-local, site-local, or organization-local.  
2763 The following special IPv6 addresses are defined:  
2764 Unspecified  
2765 An address that is not assigned to any interface and is used to indicate the absence of an  
2766 address.  
2767 Loopback  
2768 A unicast address that is not assigned to any interface and can be used by a node to send  
2769 packets to itself.  
2770 Two sets of IPv6 addresses are defined to correspond to IPv4 addresses:  
2771 IPv4-compatible addresses  
2772 These are assigned to nodes that support IPv6 and can be used when traffic is “tunneled”  
2773 through IPv4.  
2774 IPv4-mapped addresses  
2775 These are used to represent IPv4 addresses in IPv6 address format; see Section 2.10.20.2.  
2776 Note that the unspecified address and the loopback address must not be treated as IPv4-  
2777 compatible addresses.

#### 2778 2.10.20.2 Compatibility with IPv4

2779 The API provides the ability for IPv6 applications to interoperate with applications using IPv4,  
2780 by using IPv4-mapped IPv6 addresses. These addresses can be generated automatically by the  
2781 *getaddrinfo()* function when the specified host has only IPv4 addresses.

2782 Applications may use AF\_INET6 sockets to open TCP connections to IPv4 nodes, or send UDP  
2783 packets to IPv4 nodes, by simply encoding the destination's IPv4 address as an IPv4-mapped  
2784 IPv6 address, and passing that address, within a **sockaddr\_in6** structure, in the *connect()*,  
2785 *sendto()* or *sendmsg()* function. When applications use AF\_INET6 sockets to accept TCP  
2786 connections from IPv4 nodes, or receive UDP packets from IPv4 nodes, the system shall return  
2787 the peer's address to the application in the *accept()*, *recvfrom()*, *recvmsg()*, or *getpeername()*  
2788 function using a **sockaddr\_in6** structure encoded this way. If a node has an IPv4 address, then  
2789 the implementation may allow applications to communicate using that address via an  
2790 AF\_INET6 socket. In such a case, the address will be represented at the API by the  
2791 corresponding IPv4-mapped IPv6 address. Also, the implementation may allow an AF\_INET6  
2792 socket bound to **in6addr\_any** to receive inbound connections and packets destined to one of the  
2793 node's IPv4 addresses.

2794 An application may use AF\_INET6 sockets to bind to a node's IPv4 address by specifying the  
2795 address as an IPv4-mapped IPv6 address in a **sockaddr\_in6** structure in the *bind()* function. For  
2796 an AF\_INET6 socket bound to a node's IPv4 address, the system shall return the address in the  
2797 *getsockname()* function as an IPv4-mapped IPv6 address in a **sockaddr\_in6** structure.

#### 2798 2.10.20.3 Interface Identification

2799 Each local interface is assigned a unique positive integer as a numeric index. Indexes start at 1;  
2800 zero is not used. There may be gaps so that there is no current interface for a particular positive  
2801 index. Each interface also has a unique implementation-defined name.

## 2802 2.10.20.4 Options

2803 The following options apply at the IPPROTO\_IPV6 level:

## 2804 IPV6\_JOIN\_GROUP

2805 When set via *setsockopt()*, it joins the application to a multicast group on an interface  
 2806 (identified by its index) and addressed by a given multicast address, enabling packets sent  
 2807 to that address to be read via the socket. If the interface index is specified as zero, the  
 2808 system selects the interface (for example, by looking up the address in a routing table and  
 2809 using the resulting interface).

2810 An attempt to read this option using *getsockopt()* shall result in an [EOPNOTSUPP] error.

2811 The value of this option is an **ipv6\_mreq** structure.

## 2812 IPV6\_LEAVE\_GROUP

2813 When set via *setsockopt()*, it removes the application from the multicast group on an  
 2814 interface (identified by its index) and addressed by a given multicast address.

2815 An attempt to read this option using *getsockopt()* shall result in an [EOPNOTSUPP] error.

2816 The value of this option is an **ipv6\_mreq** structure.

## 2817 IPV6\_MULTICAST\_HOPS

2818 The value of this option is the hop limit for outgoing multicast IPv6 packets sent via the  
 2819 socket. Its possible values are the same as those of IPV6\_UNICAST\_HOPS. If the  
 2820 IPV6\_MULTICAST\_HOPS option is not set, a value of 1 is assumed. This option can be set  
 2821 via *setsockopt()* and read via *getsockopt()*.

## 2822 IPV6\_MULTICAST\_IF

2823 The index of the interface to be used for outgoing multicast packets. It can be set via  
 2824 *setsockopt()* and read via *getsockopt()*.

## 2825 IPV6\_MULTICAST\_LOOP

2826 This option controls whether outgoing multicast packets should be delivered back to the  
 2827 local application when the sending interface is itself a member of the destination multicast  
 2828 group. If it is set to 1 they are delivered. If it is set to 0 they are not. Other values result in an  
 2829 [EINVAL] error. This option can be set via *setsockopt()* and read via *getsockopt()*.

## 2830 IPV6\_UNICAST\_HOPS

2831 The value of this option is the hop limit for outgoing unicast IPv6 packets sent via the  
 2832 socket. If the option is not set, or is set to -1, the system selects a default value. Attempts to  
 2833 set a value less than -1 or greater than 255 shall result in an [EINVAL] error. This option can  
 2834 be set via *setsockopt()* and read via *getsockopt()*.

## 2835 IPV6\_V6ONLY

2836 This socket option restricts AF\_INET6 sockets to IPv6 communications only. AF\_INET6  
 2837 sockets may be used for both IPv4 and IPv6 communications. Some applications may want  
 2838 to restrict their use of an AF\_INET6 socket to IPv6 communications only. For these  
 2839 applications, the IPV6\_V6ONLY socket option is defined. When this option is turned on, the  
 2840 socket can be used to send and receive IPv6 packets only. This is an IPPROTO\_IPV6-level  
 2841 option. This option takes an **int** value. This is a Boolean option. By default, this option is  
 2842 turned off.

2843 An [EOPNOTSUPP] error shall result if IPV6\_JOIN\_GROUP or IPV6\_LEAVE\_GROUP is used  
 2844 with *getsockopt()*.

## 2845 2.10.20.5 Headers

2846 The symbolic constant AF\_INET6 is defined in the `<sys/socket.h>` header to identify the IPv6  
 2847 Internet address family. See the Base Definitions volume of IEEE Std 1003.1-200x, Chapter 13,  
 2848 Headers.

2849 The `sockaddr_storage` structure defined in `<sys/socket.h>` shall be large enough to  
 2850 accommodate a `sockaddr_in6` structure (see the `<netinet/in.h>` header defined in the Base  
 2851 Definitions volume of IEEE Std 1003.1-200x, Chapter 13, Headers) and shall be aligned at an  
 2852 appropriate boundary so that pointers to it can be cast as pointers to `sockaddr_in6` structures  
 2853 and used to access the fields of those structures without alignment problems. When a  
 2854 `sockaddr_storage` structure is cast as a `sockaddr_in6` structure, the `ss_family` field maps onto the  
 2855 `sin6_family` field.

2856 The `<netinet/in.h>`, `<arpa/inet.h>`, and `<netdb.h>` headers contain other definitions used in  
 2857 connection with IPv6 Internet sockets; see the Base Definitions volume of IEEE Std 1003.1-200x,  
 2858 Chapter 13, Headers.

## 2859 2.11 Tracing

2860 TRC This section describes extensions to support tracing of user applications. This functionality is  
 2861 dependent on support of the Trace option (and the rest of this section is not further shaded for  
 2862 this option).

2863 The tracing facilities defined in IEEE Std 1003.1-200x allow a process to select a set of trace event  
 2864 types, to activate a trace stream of the selected trace events as they occur in the flow of  
 2865 execution, and to retrieve the recorded trace events.

2866 The tracing operation relies on three logically different components: the traced process, the  
 2867 controller process, and the analyzer process. During the execution of the traced process, when a  
 2868 trace point is reached, a trace event is recorded into the trace streams created for that process in  
 2869 which the associated trace event type identifier is not being filtered out. The controller process  
 2870 controls the operation of recording the trace events into the trace stream. It shall be able to:

- 2871 • Initialize the attributes of a trace stream
- 2872 • Create the trace stream (for a specified traced process) using those attributes
- 2873 • Start and stop tracing for the trace stream
- 2874 • Filter the type of trace events to be recorded, if the Trace Event Filter option is supported
- 2875 • Shut a trace stream down

2876 These operations can be done for an active trace stream. The analyzer process retrieves the  
 2877 traced events either at runtime, when the trace stream has not yet been shut down, but is still  
 2878 recording trace events; or after opening a trace log that had been previously recorded and shut  
 2879 down. These three logically different operations can be performed by the same process, or can be  
 2880 distributed into different processes.

2881 A trace stream identifier can be created by a call to `posix_trace_create()`,  
 2882 `posix_trace_create_withlog()`, or `posix_trace_open()`. The `posix_trace_create()` and  
 2883 `posix_trace_create_withlog()` functions should be used by a controller process. The  
 2884 `posix_trace_open()` should be used by an analyzer process.

2885 The tracing functions can serve different purposes. One purpose is debugging the possibly pre-  
 2886 instrumented code, while another is post-mortem fault analysis. These two potential uses differ  
 2887 in that the first requires pre-filtering capabilities to avoid overwhelming the trace stream and

2888 permits focusing on expected information; while the second needs comprehensive trace  
2889 capabilities in order to be able to record all types of information.

2890 The events to be traced belong to two classes:

- 2891 1. User trace events (generated by the application instrumentation)
- 2892 2. System trace events (generated by the operating system)

2893 The trace interface defines several system trace event types associated with control of and  
2894 operation of the trace stream. This small set of system trace events includes the minimum  
2895 required to interpret correctly the trace event information present in the stream. Other desirable  
2896 system trace events for some particular application profile may be implemented and are  
2897 encouraged; for example, process and thread scheduling, signal occurrence, and so on.

2898 Each traced process shall have a mapping of the trace event names to trace event type identifiers  
2899 that have been defined for that process. Each active trace stream shall have a mapping that  
2900 incorporates all the trace event type identifiers predefined by the trace system plus all the  
2901 mappings of trace event names to trace event type identifiers of the processes that are being  
2902 traced into that trace stream. These mappings are defined from the instrumented application by  
2903 calling the *posix\_trace\_eventid\_open()* function and from the controller process by calling the  
2904 *posix\_trace\_trid\_eventid\_open()* function. For a pre-recorded trace stream, the list of trace event  
2905 types is obtained from the pre-recorded trace log.

2906 The *st\_ctime* and *st\_mtime* fields of a file associated with an active trace stream shall be marked  
2907 for update every time any of the tracing operations modifies that file.

2908 The *st\_atime* field of a file associated with a trace stream shall be marked for update every time  
2909 any of the tracing operations causes data to be read from that file.

2910 Results are undefined if the application performs any operation on a file descriptor associated  
2911 with an active or pre-recorded trace stream until *posix\_trace\_shutdown()* or *posix\_trace\_close()* is  
2912 called for that trace stream.

2913 The main purpose of this option is to define a complete set of functions and concepts that allow  
2914 a conforming application to be traced from creation to termination, whatever its realtime  
2915 constraints and properties.

## 2916 2.11.1 Tracing Data Definitions

### 2917 2.11.1.1 Structures

2918 The `<trace.h>` header shall define the *posix\_trace\_status\_info* and *posix\_trace\_event\_info* structures  
2919 described below. Implementations may add extensions to these structures.

#### 2920 **posix\_trace\_status\_info Structure**

2921 To facilitate control of a trace stream, information about the current state of an active trace  
2922 stream can be obtained dynamically. This structure is returned by a call to the  
2923 *posix\_trace\_get\_status()* function.

2924 The **posix\_trace\_status\_info** structure defined in `<trace.h>` shall contain at least the following  
2925 members:

2926  
2927  
2928  
2929  
2930  
2931

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

2932  
2933

If the Trace Log option is supported in addition to the Trace option, the **posix\_trace\_status\_info** structure defined in `<trace.h>` shall contain at least the following additional members:

2934  
2935  
2936  
2937  
2938  
2939  
2940  
2941

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

2942  
2943

The *posix\_stream\_status* member indicates the operating mode of the trace stream and shall have one of the following values defined by manifest constants in the `<trace.h>` header:

2944

POSIX\_TRACE\_RUNNING

2945

Tracing is in progress; that is, the trace stream is accepting trace events.

2946

POSIX\_TRACE\_SUSPENDED

2947

The trace stream is not accepting trace events. The tracing operation has not yet started or has stopped, either following a *posix\_trace\_stop()* function call or because the trace resources are exhausted.

2948

2949

2950

The *posix\_stream\_full\_status* member indicates the full status of the trace stream, and it shall have one of the following values defined by manifest constants in the `<trace.h>` header:

2951

2952

POSIX\_TRACE\_FULL

2953

The space in the trace stream for trace events is exhausted.

2954

POSIX\_TRACE\_NOT\_FULL

2955

There is still space available in the trace stream.

2956

The combination of the *posix\_stream\_status* and *posix\_stream\_full\_status* members also indicates the actual status of the stream. The status shall be interpreted as follows:

2957

2958

POSIX\_TRACE\_RUNNING and POSIX\_TRACE\_NOT\_FULL

2959

This status combination indicates that tracing is in progress, and there is space available for recording more trace events.

2960

2961

POSIX\_TRACE\_RUNNING and POSIX\_TRACE\_FULL

2962

This status combination indicates that tracing is in progress and that the trace stream is full of trace events. This status combination cannot occur unless the *stream-full-policy* is set to POSIX\_TRACE\_LOOP. The trace stream contains trace events recorded during a moving time window of prior trace events, and some older trace events may have been overwritten and thus lost.

2963

2964

2965

2966

2967

POSIX\_TRACE\_SUSPENDED and POSIX\_TRACE\_NOT\_FULL

2968

This status combination indicates that tracing has not yet been started, has been stopped by the *posix\_trace\_stop()* function, or has been cleared by the *posix\_trace\_clear()* function.

2969



2970 POSIX\_TRACE\_SUSPENDED and POSIX\_TRACE\_FULL  
 2971 This status combination indicates that tracing has been stopped by the implementation  
 2972 because the *stream-full-policy* attribute was POSIX\_TRACE\_UNTIL\_FULL and trace  
 2973 resources were exhausted, or that the trace stream was stopped by the function  
 2974 *posix\_trace\_stop()* at a time when trace resources were exhausted.

2975 The *posix\_stream\_outrun\_status* member indicates whether trace events were lost in the trace  
 2976 stream, and shall have one of the following values defined by manifest constants in the  
 2977 `<trace.h>` header:

2978 POSIX\_TRACE\_OVERRUN  
 2979 At least one trace event was lost and thus was not recorded in the trace stream.

2980 POSIX\_TRACE\_NO\_OVERRUN  
 2981 No trace events were lost.

2982 When the corresponding trace stream is created, the *posix\_stream\_outrun\_status* member shall be  
 2983 set to POSIX\_TRACE\_NO\_OVERRUN.

2984 Whenever an overrun occurs, *posix\_stream\_outrun\_status* member shall be set to  
 2985 POSIX\_TRACE\_OVERRUN.

2986 An overrun occurs when:

- 2987 • The policy is POSIX\_TRACE\_LOOP and a recorded trace event is overwritten.
- 2988 • The policy is POSIX\_TRACE\_UNTIL\_FULL and the trace stream is full when a trace event is  
 2989 generated.
- 2990 • If the Trace Log option is supported, the policy is POSIX\_TRACE\_FLUSH and at least one  
 2991 trace event is lost while flushing the trace stream to the trace log.

2992 The *posix\_stream\_outrun\_status* member is reset to zero after its value is read.

2993 If the Trace Log option is supported in addition to the Trace option, the *posix\_stream\_flush\_status*,  
 2994 *posix\_stream\_flush\_error*, *posix\_log\_outrun\_status*, and *posix\_log\_full\_status* members are defined  
 2995 as follows; otherwise, they are undefined.

2996 The *posix\_stream\_flush\_status* member indicates whether a flush operation is being performed  
 2997 and shall have one of the following values defined by manifest constants in the header  
 2998 `<trace.h>`:

2999 POSIX\_TRACE\_FLUSHING  
 3000 The trace stream is currently being flushed to the trace log.

3001 POSIX\_TRACE\_NOT\_FLUSHING  
 3002 No flush operation is in progress.

3003 The *posix\_stream\_flush\_status* member shall be set to POSIX\_TRACE\_FLUSHING if a flush  
 3004 operation is in progress either due to a call to the *posix\_trace\_flush()* function (explicit or caused  
 3005 by a trace stream shutdown operation) or because the trace stream has become full with the  
 3006 *stream-full-policy* attribute set to POSIX\_TRACE\_FLUSH. The *posix\_stream\_flush\_status* member  
 3007 shall be set to POSIX\_TRACE\_NOT\_FLUSHING if no flush operation is in progress.

3008 The *posix\_stream\_flush\_error* member shall be set to zero if no error occurred during flushing. If  
 3009 an error occurred during a previous flushing operation, the *posix\_stream\_flush\_error* member  
 3010 shall be set to the value of the first error that occurred. If more than one error occurs while  
 3011 flushing, error values after the first shall be discarded. The *posix\_stream\_flush\_error* member is  
 3012 reset to zero after its value is read.

3013 The *posix\_log\_outrun\_status* member indicates whether trace events were lost in the trace log,  
 3014 and shall have one of the following values defined by manifest constants in the `<trace.h>`  
 3015 header:

3016 POSIX\_TRACE\_OVERRUN  
 3017 At least one trace event was lost.

3018 POSIX\_TRACE\_NO\_OVERRUN  
 3019 No trace events were lost.

3020 When the corresponding trace stream is created, the *posix\_log\_outrun\_status* member shall be set  
 3021 to POSIX\_TRACE\_NO\_OVERRUN. Whenever an overrun occurs, this status shall be set to  
 3022 POSIX\_TRACE\_OVERRUN. The *posix\_log\_outrun\_status* member is reset to zero after its value  
 3023 is read.

3024 The *posix\_log\_full\_status* member indicates the full status of the trace log, and it shall have one of  
 3025 the following values defined by manifest constants in the `<trace.h>` header:

3026 POSIX\_TRACE\_FULL  
 3027 The space in the trace log is exhausted.

3028 POSIX\_TRACE\_NOT\_FULL  
 3029 There is still space available in the trace log.

3030 The *posix\_log\_full\_status* member is only meaningful if the *log-full-policy* attribute is either  
 3031 POSIX\_TRACE\_UNTIL\_FULL or POSIX\_TRACE\_LOOP.

3032 For an active trace stream without log, that is created by the *posix\_trace\_create()* function, the  
 3033 *posix\_log\_outrun\_status* member shall be set to POSIX\_TRACE\_NO\_OVERRUN and the  
 3034 *posix\_log\_full\_status* member shall be set to POSIX\_TRACE\_NOT\_FULL.

### 3035 **posix\_trace\_event\_info** Structure

3036 The trace event structure **posix\_trace\_event\_info** contains the information for one recorded  
 3037 trace event. This structure is returned by the set of functions *posix\_trace\_getnext\_event()*,  
 3038 *posix\_trace\_timedgetnext\_event()*, and *posix\_trace\_trygetnext\_event()*.

3039 The **posix\_trace\_event\_info** structure defined in `<trace.h>` shall contain at least the following  
 3040 members:

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

3050 In addition, if the Trace option and the Threads option are both supported, the  
 3051 **posix\_trace\_event\_info** structure defined in `<trace.h>` shall contain the following additional  
 3052 member:

3053  
3054  
3055  
3056  
3057

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

3058  
3059  
3060  
3061  
3062

The *posix\_event\_id* member represents the identification of the trace event type and its value is not directly defined by the user. This identification is returned by a call to one of the following functions: *posix\_trace\_trid\_eventid\_open()*, *posix\_trace\_eventtypelist\_getnext\_id()*, or *posix\_trace\_eventid\_open()*. The name of the trace event type can be obtained by calling *posix\_trace\_eventid\_get\_name()*.

3063  
3064  
3065

The *posix\_pid* is the process identifier of the traced process which generated the trace event. If the *posix\_event\_id* member is one of the implementation-defined system trace events and that trace event is not associated with any process, the *posix\_pid* member shall be set to zero.

3066  
3067  
3068  
3069  
3070

For a user trace event, the *posix\_prog\_address* member is the process mapped address of the point at which the associated call to the *posix\_trace\_event()* function was made. For a system trace event, if the trace event is caused by a system service explicitly called by the application, the *posix\_prog\_address* member shall be the address of the process at the point where the call to that system service was made.

3071  
3072  
3073  
3074  
3075

The *posix\_truncation\_status* member defines whether the data associated with a trace event has been truncated at the time the trace event was generated, or at the time the trace event was read from the trace stream, or (if the Trace Log option is supported) from the trace log (see the *event* argument from the *posix\_trace\_getnext\_event()* function). The *posix\_truncation\_status* member shall have one of the following values defined by manifest constants in the `<trace.h>` header:

3076  
3077

POSIX\_TRACE\_NOT\_TRUNCATED

All the traced data is available.

3078  
3079

POSIX\_TRACE\_TRUNCATED\_RECORD

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

3080  
3081  
3082  
3083

POSIX\_TRACE\_TRUNCATED\_READ

Data was truncated at the time the trace event was read from a trace stream or a trace log because the reader's buffer was too small. This truncation status overrides the POSIX\_TRACE\_TRUNCATED\_RECORD status.

3084  
3085  
3086

The *posix\_timestamp* member shall be the time at which the trace event was generated. The clock used is implementation-defined, but the resolution of this clock can be retrieved by a call to the *posix\_trace\_attr\_getclockres()* function.

3087

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

3088  
3089  
3090

- The *posix\_thread\_id* member is the identifier of the thread that generated the trace event. If the *posix\_event\_id* member is one of the implementation-defined system trace events and that trace event is not associated with any thread, the *posix\_thread\_id* member shall be set to zero.

3091

Otherwise, this member is undefined.

### 3092 2.11.1.2 Trace Stream Attributes

3093  
3094

Trace streams have attributes that compose the **posix\_trace\_attr\_t** trace stream attributes object. This object shall contain at least the following attributes:

3095

- The *generation-version* attribute identifies the origin and version of the trace system.

- 3096 • The *trace-name* attribute is a character string defined by the trace controller, and that  
3097 identifies the trace stream.
- 3098 • The *creation-time* attribute represents the time of the creation of the trace stream.
- 3099 • The *clock-resolution* attribute defines the clock resolution of the clock used to generate  
3100 timestamps.
- 3101 • The *stream-min-size* attribute defines the minimum size in bytes of the trace stream strictly  
3102 reserved for the trace events.
- 3103 • The *stream-full-policy* attribute defines the policy followed when the trace stream is full; its  
3104 value is `POSIX_TRACE_LOOP`, `POSIX_TRACE_UNTIL_FULL`, or `POSIX_TRACE_FLUSH`.
- 3105 • The *max-data-size* attribute defines the maximum record size in bytes of a trace event.

3106 In addition, if the Trace option and the Trace Inherit option are both supported, the  
3107 **posix\_trace\_attr\_t** trace stream creation attributes object shall contain at least the following  
3108 attributes:

- 3109 • The *inheritance* attribute specifies whether a newly created trace stream will inherit tracing in  
3110 its parent's process trace stream. It is either `POSIX_TRACE_INHERITED` or  
3111 `POSIX_TRACE_CLOSE_FOR_CHILD`.

3112 In addition, if the Trace option and the Trace Log option are both supported, the  
3113 **posix\_trace\_attr\_t** trace stream creation attributes object shall contain at least the following  
3114 attribute:

- 3115 • If the file type corresponding to the trace log supports the `POSIX_TRACE_LOOP` or the  
3116 `POSIX_TRACE_UNTIL_FULL` policies, the *log-max-size* attribute defines the maximum size  
3117 in bytes of the trace log associated with an active trace stream. Other stream data—for  
3118 example, trace attribute values—shall not be included in this size.
- 3119 • The *log-full-policy* attribute defines the policy of a trace log associated with an active trace  
3120 stream to be `POSIX_TRACE_LOOP`, `POSIX_TRACE_UNTIL_FULL`, or  
3121 `POSIX_TRACE_APPEND`.

## 3122 2.11.2 Trace Event Type Definitions

### 3123 2.11.2.1 System Trace Event Type Definitions

3124 The following system trace event types, defined in the `<trace.h>` header, track the invocation of  
3125 the trace operations:

- 3126 • `POSIX_TRACE_START` shall be associated with a trace start operation.
- 3127 • `POSIX_TRACE_STOP` shall be associated with a trace stop operation.
- 3128 • if the Trace Event Filter option is supported, `POSIX_TRACE_FILTER` shall be associated with  
3129 a trace event type filter change operation.

3130 The following system trace event types, defined in the `<trace.h>` header, report operational trace  
3131 events:

- 3132 • `POSIX_TRACE_OVERFLOW` shall mark the beginning of a trace overflow condition.
- 3133 • `POSIX_TRACE_RESUME` shall mark the end of a trace overflow condition.
- 3134 • If the Trace Log option is supported, `POSIX_TRACE_FLUSH_START` shall mark the  
3135 beginning of a flush operation.

3136 • If the Trace Log option is supported, POSIX\_TRACE\_FLUSH\_STOP shall mark the end of a  
 3137 flush operation.

3138 • If an implementation-defined trace error condition is reported, it shall be marked  
 3139 POSIX\_TRACE\_ERROR.

3140 The interpretation of a trace stream or a trace log by a trace analyzer process relies on the  
 3141 information recorded for each trace event, and also on system trace events that indicate the  
 3142 invocation of trace control operations and trace system operational trace events.

3143 The POSIX\_TRACE\_START and POSIX\_TRACE\_STOP trace events specify the time windows  
 3144 during which the trace stream is running.

3145 The POSIX\_TRACE\_STOP trace event with an associated data that is equal to zero indicates  
 3146 a call of the function *posix\_trace\_stop()*.

3147 The POSIX\_TRACE\_STOP trace event with an associated data that is different from zero  
 3148 indicates an automatic stop of the trace stream (see *posix\_trace\_attr\_getstreamfullpolicy()*  
 3149 defined in the System Interfaces volume of IEEE Std 1003.1-200x).

3150 The POSIX\_TRACE\_FILTER trace event indicates that a trace event type filter value changed  
 3151 while the trace stream was running.

3152 The POSIX\_TRACE\_ERROR serves to inform the analyzer process that an implementation-  
 3153 defined internal error of the trace system occurred.

3154 The POSIX\_TRACE\_OVERFLOW trace event shall be reported with a timestamp equal to the  
 3155 timestamp of the first trace event overwritten. This is an indication that some generated trace  
 3156 events have been lost.

3157 The POSIX\_TRACE\_RESUME trace event shall be reported with a timestamp equal to the  
 3158 timestamp of the first valid trace event reported after the overflow condition ends and shall be  
 3159 reported before this first valid trace event. This is an indication that the trace system is reliably  
 3160 recording trace events after an overflow condition.

3161 Each of these trace event types shall be defined by a constant trace event name and a  
 3162 **trace\_event\_id\_t** constant; trace event data is associated with some of these trace events.

3163 If the Trace option is supported and the Trace Event Filter option and the Trace Log option are  
 3164 not supported, the following predefined system trace events in Table 2-3 shall be defined:

3165 **Table 2-3** 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.

3175 If the Trace option and the Trace Event Filter option are both supported, and if the Trace Log  
 3176 option is not supported, the following predefined system trace events in Table 2-4 (on page 528)  
 3177 shall be defined:

3178

**Table 2-4** Trace and Trace Event Filter Options: System Trace Events

3179

3180

3181

3182

3183

3184

3185

3186

3187

3188

3189

3190

3191

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.

3192

3193

3194

If the Trace option and the Trace Log option are both supported, and if the Trace Event Filter option is not supported, the following predefined system trace events in Table 2-5 shall be defined:

3195

**Table 2-5** Trace and Trace Log Options: System Trace Events

3196

3197

3198

3199

3200

3201

3202

3203

3204

3205

3206

3207

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.

3208

3209

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 2-6 (on page 529) shall be defined:

3210 **Table 2-6** Trace, Trace Log, 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.
posix_trace_flush_start	POSIX_TRACE_FLUSH_START	None.
posix_trace_flush_stop	POSIX_TRACE_FLUSH_STOP	None.

3226 **2.11.2.2** *User Trace Event Type Definitions*

3227 The user trace event `POSIX_TRACE_UNNAMED_USEREVENT` is defined in the `<trace.h>`  
 3228 header. If the limit of per-process user trace event names represented by  
 3229 `{TRACE_USER_EVENT_MAX}` has already been reached, this predefined user event shall be  
 3230 returned when the application tries to register more events than allowed. The data associated  
 3231 with this trace event is application-defined.

3232 The following predefined user trace event in Table 2-7 shall be defined:

3233 **Table 2-7** Trace Option: User Trace Event

Event Name	Constant
posix_trace_unnamed_userevent	POSIX_TRACE_UNNAMED_USEREVENT

3236 **2.11.3** **Trace Functions**

3237 The trace interface is built and structured to improve portability through use of trace data of  
 3238 opaque type. The object-oriented approach for the manipulation of trace attributes and trace  
 3239 event type identifiers requires definition of many constructor and selector functions which  
 3240 operate on these opaque types. Also, the trace interface must support several different tracing  
 3241 roles. To facilitate reading the trace interface, the trace functions are grouped into small  
 3242 functional sets supporting the three different roles:

- 3243 • A trace controller process requires functions to set up and customize all the resources needed  
 3244 to run a trace stream, including:
  - 3245 — Attribute initialization and destruction (`posix_trace_attr_init()`)
  - 3246 — Identification information manipulation (`posix_trace_attr_getgenversion()`)
  - 3247 — Trace system behavior modification (`posix_trace_attr_getinherited()`)
  - 3248 — Trace stream and trace log size set (`posix_trace_attr_getmaxusereventsize()`)

- 3249 — Trace stream creation, flush, and shutdown (*posix\_trace\_create()*)
- 3250 — Trace stream and trace log clear (*posix\_trace\_clear()*)
- 3251 — Trace event type identifier manipulation (*posix\_trace\_trid\_eventid\_open()*)
- 3252 — Trace event type identifier list exploration (*posix\_trace\_eventtypelist\_getnext\_id()*)
- 3253 — Trace event type set manipulation (*posix\_trace\_eventset\_empty()*)
- 3254 — Trace event type filter set (*posix\_trace\_set\_filter()*)
- 3255 — Trace stream start and stop (*posix\_trace\_start()*)
- 3256 — Trace stream information and status read (*posix\_trace\_get\_attr()*)
- 3257 • A traced process requires functions to instrument trace points:
- 3258 — Trace event type identifiers definition and trace points insertion (*posix\_trace\_event()*)
- 3259 • A trace analyzer process requires functions to retrieve information from a trace stream and
- 3260 trace log:
- 3261 — Identification information read (*posix\_trace\_attr\_getgenversion()*)
- 3262 — Trace system behavior information read (*posix\_trace\_attr\_getinherited()*)
- 3263 — Trace stream and trace log size get (*posix\_trace\_attr\_getmaxusereventsized()*)
- 3264 — Trace event type identifier manipulation (*posix\_trace\_trid\_eventid\_open()*)
- 3265 — Trace event type identifier list exploration (*posix\_trace\_eventtypelist\_getnext\_id()*)
- 3266 — Trace log open, rewind, and close (*posix\_trace\_open()*)
- 3267 — Trace stream information and status read (*posix\_trace\_get\_attr()*)
- 3268 — trace event read (*posix\_trace\_getnext\_event()*)

## 3269 2.12 Data Types

3270 All of the data types used by various functions are defined by the implementation. The  
 3271 following table describes some of these types. Other types referenced in the description of a  
 3272 function, not mentioned here, can be found in the appropriate header for that function.

3273

3274

3275

3276

3277

3278

3279

3280

3281

3282

3283

3284

Defined Type	Description
<b>cc_t</b>	Type used for terminal special characters.
<b>clock_t</b>	Arithmetic type used for processor times, as defined in the ISO C standard.
<b>clockid_t</b>	Used for clock ID type in some timer functions.
<b>dev_t</b>	Arithmetic type used for device numbers.
<b>DIR</b>	Type representing a directory stream.
<b>div_t</b>	Structure type returned by the <i>div()</i> function.
<b>FILE</b>	Structure containing information about a file.
<b>glob_t</b>	Structure type used in pathname pattern matching.
<b>fpos_t</b>	Type containing all information needed to specify uniquely every



Defined Type	Description
3285	position within a file.
3286	
3287	
3288	<b>gid_t</b> Arithmetic type used for group IDs.
3289	<b>iconv_t</b> Type used for conversion descriptors.
3290	<b>id_t</b> Arithmetic type used as a general identifier; can be used to contain at least the largest of a <b>pid_t</b> , <b>uid_t</b> , or <b>gid_t</b> .
3291	
3292	<b>ino_t</b> Arithmetic type used for file serial numbers.
3293	<b>key_t</b> Arithmetic type used for XSI interprocess communication.
3294	<b>ldiv_t</b> Structure type returned by the <i>ldiv()</i> function.
3295	<b>mode_t</b> Arithmetic type used for file attributes.
3296	<b>mqd_t</b> Used for message queue descriptors.
3297	<b>nfds_t</b> Integer type used for the number of file descriptors.
3298	<b>nlink_t</b> Arithmetic type used for link counts.
3299	<b>off_t</b> Signed arithmetic type used for file sizes.
3300	<b>pid_t</b> Signed arithmetic type used for process and process group IDs.
3301	<b>pthread_attr_t</b> Used to identify a thread attribute object.
3302	<b>pthread_cond_t</b> Used for condition variables.
3303	<b>pthread_condattr_t</b> Used to identify a condition attribute object.
3304	<b>pthread_key_t</b> Used for thread-specific data keys.
3305	<b>pthread_mutex_t</b> Used for mutexes.
3306	<b>pthread_mutexattr_t</b> Used to identify a mutex attribute object.
3307	<b>pthread_once_t</b> Used for dynamic package initialization.
3308	<b>pthread_rwlock_t</b> Used for read-write locks.
3309	<b>pthread_rwlockattr_t</b> Used for read-write lock attributes.
3310	<b>pthread_t</b> Used to identify a thread.
3311	<b>ptrdiff_t</b> Signed integer type of the result of subtracting two pointers.
3312	<b>regex_t</b> Structure type used in regular expression matching.
3313	<b>regmatch_t</b> Structure type used in regular expression matching.
3314	<b>rlim_t</b> Unsigned arithmetic type used for limit values, to which objects of type <b>int</b> and <b>off_t</b> can be cast without loss of value.
3315	
3316	<b>sem_t</b> Type used in performing semaphore operations.
3317	<b>sig_atomic_t</b> Integer type of an object that can be accessed as an atomic entity, even in the presence of asynchronous interrupts.
3318	
3319	<b>sigset_t</b> Integer or structure type of an object used to represent sets of signals.
3320	
3321	<b>size_t</b> Unsigned integer type used for size of objects.
3322	<b>speed_t</b> Type used for terminal baud rates.
3323	<b>ssize_t</b> Arithmetic type used for a count of bytes or an error indication.
3324	<b>suseconds_t</b> Signed arithmetic type used for time in microseconds.
3325	<b>tcflag_t</b> Type used for terminal modes.
3326	<b>time_t</b> Arithmetic type used for time in seconds, as defined in the ISO C standard.
3327	
3328	<b>timer_t</b> Used for timer ID returned by the <i>timer_create()</i> function.
3329	<b>uid_t</b> Arithmetic type used for user IDs.
3330	<b>useconds_t</b> Integer type used for time in microseconds.
3331	<b>va_list</b> Type used for traversing variable argument lists.
3332	<b>wchar_t</b> Integer type whose range of values can represent distinct codes for all members of the largest extended character set specified by the
3333	

3334  
3335  
3336  
3337  
3338  
3339  
3340  
3341

Defined Type	Description
<b>wctype_t</b>	supported locales.
<b>wint_t</b>	Scalar type which represents a character class descriptor.
<b>wint_t</b>	Integer type capable of storing any valid value of <b>wchar_t</b> or WEOF.
<b>wordexp_t</b>	Structure type used in word expansion.

# *System Interfaces*

3342

3343  
3344

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

3345 **NAME**

3346       FD\_CLR — macros for synchronous I/O multiplexing

3347 **SYNOPSIS**

3348       #include &lt;sys/time.h&gt;

3349       FD\_CLR(int *fd*, fd\_set \**fdset*);3350       FD\_ISSET(int *fd*, fd\_set \**fdset*);3351       FD\_SET(int *fd*, fd\_set \**fdset*);3352       FD\_ZERO(fd\_set \**fdset*);3353 **DESCRIPTION**3354       Refer to *pselect*(.).

|

3355 **NAME**

3356        \_Exit, \_exit — terminate a process

3357 **SYNOPSIS**

3358        #include &lt;unistd.h&gt;

3359        void \_Exit(int *status*);3360        void \_exit(int *status*);3361 **DESCRIPTION**3362        Refer to *exit()*.

3363 **NAME**

3364        \_longjmp, \_setjmp — non-local goto

3365 **SYNOPSIS**

3366 XSI       #include <setjmp.h>

3367       void \_longjmp(jmp\_buf env, int val);

3368       int \_setjmp(jmp\_buf env);

3369

3370 **DESCRIPTION**

3371       The *\_longjmp()* and *\_setjmp()* functions shall be equivalent to *longjmp()* and *setjmp()*,  
3372       respectively, with the additional restriction that *\_longjmp()* and *\_setjmp()* shall not manipulate  
3373       the signal mask.

3374       If *\_longjmp()* is called even though *env* was never initialized by a call to *\_setjmp()*, or when the  
3375       last such call was in a function that has since returned, the results are undefined.

3376 **RETURN VALUE**

3377       Refer to *longjmp()* and *setjmp()*.

3378 **ERRORS**

3379       No errors are defined.

3380 **EXAMPLES**

3381       None.

3382 **APPLICATION USAGE**

3383       If *\_longjmp()* is executed and the environment in which *\_setjmp()* was executed no longer exists,  
3384       errors can occur. The conditions under which the environment of the *\_setjmp()* no longer exists  
3385       include exiting the function that contains the *\_setjmp()* call, and exiting an inner block with  
3386       temporary storage. This condition might not be detectable, in which case the *\_longjmp()* occurs  
3387       and, if the environment no longer exists, the contents of the temporary storage of an inner block  
3388       are unpredictable. This condition might also cause unexpected process termination. If the  
3389       function has returned, the results are undefined.

3390       Passing *longjmp()* a pointer to a buffer not created by *setjmp()*, passing *\_longjmp()* a pointer to a  
3391       buffer not created by *\_setjmp()*, passing *siglongjmp()* a pointer to a buffer not created by  
3392       *sigsetjmp()*, or passing any of these three functions a buffer that has been modified by the user  
3393       can cause all the problems listed above, and more.

3394       The *\_longjmp()* and *\_setjmp()* functions are included to support programs written to historical  
3395       system interfaces. New applications should use *siglongjmp()* and *sigsetjmp()* respectively.

3396 **RATIONALE**

3397       None.

3398 **FUTURE DIRECTIONS**

3399       The *\_longjmp()* and *\_setjmp()* functions may be marked LEGACY in a future version.

3400 **SEE ALSO**

3401       *longjmp()*, *setjmp()*, *siglongjmp()*, *sigsetjmp()*, the Base Definitions volume of  
3402       IEEE Std 1003.1-200x, <setjmp.h>

3403 **CHANGE HISTORY**

3404       First released in Issue 4, Version 2.

3405 **Issue 5**

3406 Moved from X/OPEN UNIX extension to BASE.

3407 **NAME**

3408        \_setjmp — set jump point for a non-local goto

3409 **SYNOPSIS**

3410 xSI     #include <setjmp.h>

3411        int \_setjmp( jmp\_buf env );

3412

3413 **DESCRIPTION**

3414        Refer to *\_longjmp()*.



3415 **NAME**

3416        \_toupper — transliterate uppercase characters to lowercase

3417 **SYNOPSIS**

3418 xSI        #include &lt;ctype.h&gt;

3419        int \_tolower(int c);

3420

3421 **DESCRIPTION**3422        The *\_tolower()* macro shall be equivalent to *tolower(c)* except that the application shall ensure  
3423        that the argument *c* is an uppercase letter.3424 **RETURN VALUE**3425        Upon successful completion, *\_tolower()* shall return the lowercase letter corresponding to the  
3426        argument passed.3427 **ERRORS**

3428        No errors are defined.

3429 **EXAMPLES**

3430        None.

3431 **APPLICATION USAGE**

3432        None.

3433 **RATIONALE**

3434        None.

3435 **FUTURE DIRECTIONS**

3436        None.

3437 **SEE ALSO**3438        *tolower()*, *isupper()*, the Base Definitions volume of IEEE Std 1003.1-200x, <ctype.h>, the Base  
3439        Definitions volume of IEEE Std 1003.1-200x, Chapter 7, Locale3440 **CHANGE HISTORY**

3441        First released in Issue 1. Derived from Issue 1 of the SVID.

3442 **Issue 6**

3443        The DESCRIPTION is updated to avoid use of the term “must” for application requirements.

3444 **NAME**

3445        \_toupper — transliterate lowercase characters to uppercase

3446 **SYNOPSIS**

```
3447 xSI     #include <ctype.h>
```

```
3448        int _toupper(int c);
```

3449

3450 **DESCRIPTION**

3451        The *\_toupper()* macro shall be equivalent to *toupper()* except that the application shall ensure  
3452        that the argument *c* is a lowercase letter.

3453 **RETURN VALUE**

3454        Upon successful completion, *\_toupper()* shall return the uppercase letter corresponding to the  
3455        argument passed.

3456 **ERRORS**

3457        No errors are defined.

3458 **EXAMPLES**

3459        None.

3460 **APPLICATION USAGE**

3461        None.

3462 **RATIONALE**

3463        None.

3464 **FUTURE DIRECTIONS**

3465        None.

3466 **SEE ALSO**

3467        *islower()*, *toupper()*, the Base Definitions volume of IEEE Std 1003.1-200x, **<ctype.h>**, the Base  
3468        Definitions volume of IEEE Std 1003.1-200x, Chapter 7, Locale

3469 **CHANGE HISTORY**

3470        First released in Issue 1. Derived from Issue 1 of the SVID.

3471 **Issue 6**

3472        The DESCRIPTION is updated to avoid use of the term “must” for application requirements.

3473 **NAME**

3474 a64l, l64a — convert between a 32-bit integer and a radix-64 ASCII string

3475 **SYNOPSIS**

```
3476 xSI #include <stdlib.h>
3477 long a64l(const char *s);
3478 char *l64a(long value);
3479
```

3480 **DESCRIPTION**

3481 These functions maintain numbers stored in radix-64 ASCII characters. This is a notation by  
 3482 which 32-bit integers can be represented by up to six characters; each character represents a digit  
 3483 in radix-64 notation. If the type **long** contains more than 32 bits, only the low-order 32 bits shall  
 3484 be used for these operations.

3485 The characters used to represent digits are '.' (dot) for 0, '/' for 1, '0' through '9' for [2,11],  
 3486 'A' through 'Z' for [12,37], and 'a' through 'z' for [38,63].

3487 The *a64l()* function shall take a pointer to a radix-64 representation, in which the first digit is the  
 3488 least significant, and return the corresponding **long** value. If the string pointed to by *s* contains  
 3489 more than six characters, *a64l()* shall use the first six. If the first six characters of the string  
 3490 contain a null terminator, *a64l()* shall use only characters preceding the null terminator. The  
 3491 *a64l()* function shall scan the character string from left to right with the least significant digit on  
 3492 the left, decoding each character as a 6-bit radix-64 number. If the type **long** contains more than  
 3493 32 bits, the resulting value is sign-extended. The behavior of *a64l()* is unspecified if *s* is a null  
 3494 pointer or the string pointed to by *s* was not generated by a previous call to *l64a()*.

3495 The *l64a()* function shall take a **long** argument and return a pointer to the corresponding radix-  
 3496 64 representation. The behavior of *l64a()* is unspecified if *value* is negative.

3497 The value returned by *l64a()* may be a pointer into a static buffer. Subsequent calls to *l64a()* may  
 3498 overwrite the buffer.

3499 The *l64a()* function need not be reentrant. A function that is not required to be reentrant is not  
 3500 required to be thread-safe.

3501 **RETURN VALUE**

3502 Upon successful completion, *a64l()* shall return the **long** value resulting from conversion of the  
 3503 input string. If a string pointed to by *s* is an empty string, *a64l()* shall return 0L.

3504 The *l64a()* function shall return a pointer to the radix-64 representation. If *value* is 0L, *l64a()* shall  
 3505 return a pointer to an empty string.

3506 **ERRORS**

3507 No errors are defined.

3508 **EXAMPLES**

3509 None.

3510 **APPLICATION USAGE**

3511 If the type **long** contains more than 32 bits, the result of *a64l(l64a(x))* is *x* in the low-order 32 bits.

3512 **RATIONALE**

3513 This is not the same encoding as used by either encoding variant of the *uuencode* utility.

3514 **FUTURE DIRECTIONS**

3515 None.

3516 **SEE ALSO**3517 *strtoul()*, the Base Definitions volume of IEEE Std 1003.1-200x, `<stdlib.h>`, the Shell and Utilities  
3518 volume of IEEE Std 1003.1-200x, *uuencode*3519 **CHANGE HISTORY**

3520 First released in Issue 4, Version 2.

3521 **Issue 5**

3522 Moved from X/OPEN UNIX extension to BASE.

3523 Normative text previously in the APPLICATION USAGE section moved to the DESCRIPTION.

3524 A note indicating that these functions need not be reentrant is added to the DESCRIPTION.

3525 **NAME**

3526 abort — generate an abnormal process abort

3527 **SYNOPSIS**

3528 #include &lt;stdlib.h&gt;

3529 void abort(void);

3530 **DESCRIPTION**

3531 CX The functionality described on this reference page is aligned with the ISO C standard. Any  
 3532 conflict between the requirements described here and the ISO C standard is unintentional. This  
 3533 volume of IEEE Std 1003.1-200x defers to the ISO C standard.

3534 The *abort()* function shall cause abnormal process termination to occur, unless the signal  
 3535 SIGABRT is being caught and the signal handler does not return.

3536 CX The abnormal termination processing shall include at least the effect of *fclose()* on all open  
 3537 streams and the default actions defined for SIGABRT.

3538 XSI On XSI-conformant systems, in addition the abnormal termination processing shall include the  
 3539 effect of *fclose()* on message catalog descriptors.

3540 The SIGABRT signal shall be sent to the calling process as if by means of *raise()* with the  
 3541 argument SIGABRT.

3542 CX The status made available to *wait()* or *waitpid()* by *abort()* shall be that of a process terminated  
 3543 by the SIGABRT signal. The *abort()* function shall override blocking or ignoring the SIGABRT  
 3544 signal. |

3545 **RETURN VALUE**3546 The *abort()* function shall not return.3547 **ERRORS**

3548 No errors are defined.

3549 **EXAMPLES**

3550 None.

3551 **APPLICATION USAGE**

3552 Catching the signal is intended to provide the application writer with a portable means to abort  
 3553 processing, free from possible interference from any implementation-defined functions.

3554 **RATIONALE**

3555 None.

3556 **FUTURE DIRECTIONS**

3557 None.

3558 **SEE ALSO**

3559 *exit()*, *kill()*, *raise()*, *signal()*, *wait()*, *waitpid()*, the Base Definitions volume of  
 3560 IEEE Std 1003.1-200x, <stdlib.h>

3561 **CHANGE HISTORY**

3562 First released in Issue 1. Derived from Issue 1 of the SVID.

3563 **Issue 6**

3564 Extensions beyond the ISO C standard are now marked.

3565 **NAME**

3566       abs — return an integer absolute value

3567 **SYNOPSIS**

3568       #include &lt;stdlib.h&gt;

3569       int abs(int i);

3570 **DESCRIPTION**

3571 cx     The functionality described on this reference page is aligned with the ISO C standard. Any  
3572     conflict between the requirements described here and the ISO C standard is unintentional. This  
3573     volume of IEEE Std 1003.1-200x defers to the ISO C standard.

3574     The *abs()* function shall compute the absolute value of its integer operand, *i*. If the result cannot  
3575     be represented, the behavior is undefined.

3576 **RETURN VALUE**3577       The *abs()* function shall return the absolute value of its integer operand.3578 **ERRORS**

3579       No errors are defined.

3580 **EXAMPLES**

3581       None.

3582 **APPLICATION USAGE**

3583       In two's-complement representation, the absolute value of the negative integer with largest  
3584       magnitude {INT\_MIN} might not be representable.

3585 **RATIONALE**

3586       None.

3587 **FUTURE DIRECTIONS**

3588       None.

3589 **SEE ALSO**3590       *fabs()*, *labs()*, the Base Definitions volume of IEEE Std 1003.1-200x, <stdlib.h>3591 **CHANGE HISTORY**

3592       First released in Issue 1. Derived from Issue 1 of the SVID.

3593 **Issue 6**

3594       Extensions beyond the ISO C standard are now marked.

3595 **NAME**

3596            accept — accept a new connection on a socket

3597 **SYNOPSIS**

3598            #include &lt;sys/socket.h&gt;

3599            int accept(int *socket*, struct sockaddr \*restrict *address*,  
3600                socklen\_t \*restrict *address\_len*);3601 **DESCRIPTION**3602            The *accept()* function shall extract the first connection on the queue of pending connections,  
3603            create a new socket with the same socket type protocol and address family as the specified  
3604            socket, and allocate a new file descriptor for that socket.3605            The *accept()* function takes the following arguments:3606            *socket*                Specifies a socket that was created with *socket()*, has been bound to an address  
3607                                    with *bind()*, and has issued a successful call to *listen()*.3608            *address*               Either a null pointer, or a pointer to a **sockaddr** structure where the address of  
3609                                    the connecting socket shall be returned.3610            *address\_len*           Points to a **socklen\_t** structure which on input specifies the length of the  
3611                                    supplied **sockaddr** structure, and on output specifies the length of the stored  
3612                                    address.3613            If *address* is not a null pointer, the address of the peer for the accepted connection shall be stored  
3614            in the **sockaddr** structure pointed to by *address*, and the length of this address shall be stored in  
3615            the object pointed to by *address\_len*.3616            If the actual length of the address is greater than the length of the supplied **sockaddr** structure,  
3617            the stored address shall be truncated.3618            If the protocol permits connections by unbound clients, and the peer is not bound, then the value  
3619            stored in the object pointed to by *address* is unspecified.3620            If the listen queue is empty of connection requests and O\_NONBLOCK is not set on the file  
3621            descriptor for the socket, *accept()* shall block until a connection is present. If the *listen()* queue is  
3622            empty of connection requests and O\_NONBLOCK is set on the file descriptor for the socket,  
3623            *accept()* shall fail and set *errno* to [EAGAIN] or [EWOULDBLOCK].3624            The accepted socket cannot itself accept more connections. The original socket remains open and  
3625            can accept more connections.3626 **RETURN VALUE**3627            Upon successful completion, *accept()* shall return the non-negative file descriptor of the accepted  
3628            socket. Otherwise, -1 shall be returned and *errno* set to indicate the error.3629 **ERRORS**3630            The *accept()* function shall fail if:

3631            [EAGAIN] or [EWOULDBLOCK]

3632                                    O\_NONBLOCK is set for the socket file descriptor and no connections are  
3633                                    present to be accepted.3634            [EBADF]                The *socket* argument is not a valid file descriptor.

3635            [ECONNABORTED]

3636                                    A connection has been aborted.

3637	[EINTR]	The <i>accept()</i> function was interrupted by a signal that was caught before a valid connection arrived.
3638		
3639	[EINVAL]	The <i>socket</i> is not accepting connections.
3640	[EMFILE]	{OPEN_MAX} file descriptors are currently open in the calling process.
3641	[ENFILE]	The maximum number of file descriptors in the system are already open.
3642	[ENOTSOCK]	The <i>socket</i> argument does not refer to a socket.
3643	[EOPNOTSUPP]	The socket type of the specified socket does not support accepting connections.
3644		
3645		The <i>accept()</i> function may fail if:
3646	[ENOBUFS]	No buffer space is available.
3647	[ENOMEM]	There was insufficient memory available to complete the operation.
3648	XSR [EPROTO]	A protocol error has occurred; for example, the STREAMS protocol stack has not been initialized.
3649		

3650 **EXAMPLES**

3651 None.

3652 **APPLICATION USAGE**

3653 When a connection is available, *select()* indicates that the file descriptor for the socket is ready  
 3654 for reading.

3655 **RATIONALE**

3656 None.

3657 **FUTURE DIRECTIONS**

3658 None.

3659 **SEE ALSO**

3660 *bind()*, *connect()*, *listen()*, *socket()*, the Base Definitions volume of IEEE Std 1003.1-200x,  
 3661 <sys/socket.h>

3662 **CHANGE HISTORY**

3663 First released in Issue 6. Derived from the XNS, Issue 5.2 specification.

3664 The **restrict** keyword is added to the *accept()* prototype for alignment with the  
 3665 ISO/IEC 9899:1999 standard.



3666 **NAME**

3667           access — determine accessibility of a file

3668 **SYNOPSIS**

3669           #include &lt;unistd.h&gt;

3670           int access(const char \*path, int amode);

3671 **DESCRIPTION**

3672           The `access()` function shall check the file named by the pathname pointed to by the `path` |  
 3673           argument for accessibility according to the bit pattern contained in `amode`, using the real user ID |  
 3674           in place of the effective user ID and the real group ID in place of the effective group ID.

3675           The value of `amode` is either the bitwise-inclusive OR of the access permissions to be checked |  
 3676           (R\_OK, W\_OK, X\_OK) or the existence test (F\_OK).

3677           If any access permissions are checked, each shall be checked individually, as described in the |  
 3678           Base Definitions volume of IEEE Std 1003.1-200x, Chapter 3, Definitions. If the process has |  
 3679           appropriate privileges, an implementation may indicate success for X\_OK even if none of the |  
 3680           execute file permission bits are set.

3681 **RETURN VALUE**

3682           If the requested access is permitted, `access()` succeeds and shall return 0; otherwise, -1 shall be |  
 3683           returned and `errno` shall be set to indicate the error.

3684 **ERRORS**3685           The `access()` function shall fail if:

3686           [EACCES]           Permission bits of the file mode do not permit the requested access, or search |  
 3687           permission is denied on a component of the path prefix.

3688           [ELOOP]           A loop exists in symbolic links encountered during resolution of the `path` |  
 3689           argument.

3690           [ENAMETOOLONG]

3691                               The length of the `path` argument exceeds {PATH\_MAX} or a pathname |  
 3692                               component is longer than {NAME\_MAX}. |

3693           [ENOENT]           A component of `path` does not name an existing file or `path` is an empty string.

3694           [ENOTDIR]          A component of the path prefix is not a directory.

3695           [EROFS]           Write access is requested for a file on a read-only file system.

3696           The `access()` function may fail if:

3697           [EINVAL]           The value of the `amode` argument is invalid.

3698           [ELOOP]           More than {SYMLOOP\_MAX} symbolic links were encountered during |  
 3699           resolution of the `path` argument.

3700           [ENAMETOOLONG]

3701                               As a result of encountering a symbolic link in resolution of the `path` argument, |  
 3702                               the length of the substituted pathname string exceeded {PATH\_MAX}. |

3703           [ETXTBSY]          Write access is requested for a pure procedure (shared text) file that is being |  
 3704           executed.

3705 **EXAMPLES**3706 **Testing for the Existence of a File**

3707 The following example tests whether a file named **myfile** exists in the **/tmp** directory.

```
3708 #include <unistd.h>
3709 ...
3710 int result;
3711 const char *filename = "/tmp/myfile";
3712 result = access (filename, F_OK);
```

3713 **APPLICATION USAGE**

3714 Additional values of *amode* other than the set defined in the description may be valid; for  
3715 example, if a system has extended access controls.

3716 **RATIONALE**

3717 In early proposals, some inadequacies in the *access()* function led to the creation of an *eaccess()*  
3718 function because:

- 3719 1. Historical implementations of *access()* do not test file access correctly when the process'  
3720 real user ID is superuser. In particular, they always return zero when testing execute  
3721 permissions without regard to whether the file is executable.
- 3722 2. The superuser has complete access to all files on a system. As a consequence, programs  
3723 started by the superuser and switched to the effective user ID with lesser privileges cannot  
3724 use *access()* to test their file access permissions.

3725 However, the historical model of *eaccess()* does not resolve problem (1), so this volume of  
3726 IEEE Std 1003.1-200x now allows *access()* to behave in the desired way because several  
3727 implementations have corrected the problem. It was also argued that problem (2) is more easily  
3728 solved by using *open()*, *chdir()*, or one of the *exec* functions as appropriate and responding to the  
3729 error, rather than creating a new function that would not be as reliable. Therefore, *eaccess()* is not  
3730 included in this volume of IEEE Std 1003.1-200x.

3731 The sentence concerning appropriate privileges and execute permission bits reflects the two  
3732 possibilities implemented by historical implementations when checking superuser access for  
3733 X\_OK.

3734 New implementations are discouraged from returning X\_OK unless at least one execution  
3735 permission bit is set.

3736 **FUTURE DIRECTIONS**

3737 None.

3738 **SEE ALSO**

3739 *chmod()*, *stat()*, the Base Definitions volume of IEEE Std 1003.1-200x, **<unistd.h>**

3740 **CHANGE HISTORY**

3741 First released in Issue 1. Derived from Issue 1 of the SVID.

3742 **Issue 6**

3743 The following new requirements on POSIX implementations derive from alignment with the  
3744 Single UNIX Specification:

- 3745 • The [ELOOP] mandatory error condition is added.
- 3746 • A second [ENAMETOOLONG] is added as an optional error condition.

3747

- The [ETXTBSY] optional error condition is added.

3748

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

3749

- The [ELOOP] optional error condition is added.

3750 **NAME**

3751           acos, acosf, acosl — arc cosine functions

3752 **SYNOPSIS**

3753           #include &lt;math.h&gt;

3754           double acos(double x);

3755           float acosf(float x);

3756           long double acosl(long double x);

3757 **DESCRIPTION**

3758 **CX**       The functionality described on this reference page is aligned with the ISO C standard. Any  
 3759 conflict between the requirements described here and the ISO C standard is unintentional. This  
 3760 volume of IEEE Std 1003.1-200x defers to the ISO C standard.

3761       These functions shall compute the principal value of the arc cosine of their argument *x*. The  
 3762 value of *x* should be in the range  $[-1,1]$ .

3763       An application wishing to check for error situations should set *errno* to zero and call  
 3764 *feclearexcept*(*FE\_ALL\_EXCEPT*) before calling these functions. On return, if *errno* is non-zero or  
 3765 *fetestexcept*(*FE\_INVALID* | *FE\_DIVBYZERO* | *FE\_OVERFLOW* | *FE\_UNDERFLOW*) is non-  
 3766 zero, an error has occurred.

3767 **RETURN VALUE**

3768       Upon successful completion, these functions shall return the arc cosine of *x*, in the range  $[0,\pi]$   
 3769 radians.

3770 **MX**       For finite values of *x* not in the range  $[-1,1]$ , a domain error shall occur, and either a NaN (if  
 3771 supported), or an implementation-defined value shall be returned.

3772 **MX**       If *x* is NaN, a NaN shall be returned.

3773       If *x* is +1, +0 shall be returned.

3774       If *x* is  $\pm\text{Inf}$ , a domain error shall occur, and either a NaN (if supported), or an implementation-  
 3775 defined value shall be returned.

3776 **ERRORS**

3777       These functions shall fail if:

3778 **MX**       Domain Error    The *x* argument is finite and is not in the range  $[-1,1]$ , or is  $\pm\text{Inf}$ .

3779               If the integer expression (*math\_errhandling* & *MATH\_ERRNO*) is non-zero, |  
 3780 then *errno* shall be set to [EDOM]. If the integer expression (*math\_errhandling* |  
 3781 & *MATH\_ERREXCEPT*) is non-zero, then the invalid floating-point exception |  
 3782 shall be raised. |

3783 **EXAMPLES**

3784       None.

3785 **APPLICATION USAGE**

3786       On error, the expressions (*math\_errhandling* & *MATH\_ERRNO*) and (*math\_errhandling* &  
 3787 *MATH\_ERREXCEPT*) are independent of each other, but at least one of them must be non-zero.

3788 **RATIONALE**

3789       None.

3790 **FUTURE DIRECTIONS**

3791 None.

3792 **SEE ALSO**

3793 *cos()*, *feclearexcept()*, *fetetestexcept()*, *isnan()*, the Base Definitions volume of IEEE Std 1003.1-200x, |  
3794 Section 4.18, Treatment of Error Conditions for Mathematical Functions, <math.h> |

3795 **CHANGE HISTORY**

3796 First released in Issue 1. Derived from Issue 1 of the SVID.

3797 **Issue 5**

3798 The DESCRIPTION is updated to indicate how an application should check for an error. This  
3799 text was previously published in the APPLICATION USAGE section.

3800 **Issue 6**3801 The *acosf()* and *acosl()* functions are added for alignment with the ISO/IEC 9899:1999 standard.

3802 The DESCRIPTION, RETURN VALUE, ERRORS, and APPLICATION USAGE sections are  
3803 revised to align with the ISO/IEC 9899:1999 standard.

3804 IEC 60559:1989 standard floating-point extensions over the ISO/IEC 9899:1999 standard are  
3805 marked.

3806 **NAME**3807        **acosf** — arc cosine functions3808 **SYNOPSIS**

3809        #include &lt;math.h&gt;

3810        float acosf(float x);

3811 **DESCRIPTION**3812        Refer to *acos()*.

3813 **NAME**

3814 acosh, acoshf, acoshl, — inverse hyperbolic cosine functions

3815 **SYNOPSIS**

3816 #include &lt;math.h&gt;

3817 double acosh(double x);

3818 float acoshf(float x);

3819 long double acoshl(long double x);

3820 **DESCRIPTION**

3821 **CX** The functionality described on this reference page is aligned with the ISO C standard. Any  
 3822 conflict between the requirements described here and the ISO C standard is unintentional. This  
 3823 volume of IEEE Std 1003.1-200x defers to the ISO C standard.

3824 These functions shall compute the inverse hyperbolic cosine of their argument *x*.

3825 An application wishing to check for error situations should set *errno* to zero and call  
 3826 *feclearexcept*(FE\_ALL\_EXCEPT) before calling these functions. On return, if *errno* is non-zero or  
 3827 *fetestexcept*(FE\_INVALID | FE\_DIVBYZERO | FE\_OVERFLOW | FE\_UNDERFLOW) is non-  
 3828 zero, an error has occurred.

3829 **RETURN VALUE**3830 Upon successful completion, these functions shall return the inverse hyperbolic cosine of their  
3831 argument.3832 **MX** For finite values of  $x < 1$ , a domain error shall occur, and either a NaN (if supported), or an  
3833 implementation-defined value shall be returned.3834 **MX** If *x* is NaN, a NaN shall be returned.3835 If *x* is +1, +0 shall be returned.3836 If *x* is +Inf, +Inf shall be returned.3837 If *x* is -Inf, a domain error shall occur, and either a NaN (if supported), or an implementation-  
3838 defined value shall be returned.3839 **ERRORS**

3840 These functions shall fail if:

3841 **MX** Domain Error The *x* argument is finite and less than +1.0, or is -Inf.

3842 If the integer expression (math\_errhandling & MATH\_ERRNO) is non-zero, |  
 3843 then *errno* shall be set to [EDOM]. If the integer expression (math\_errhandling |  
 3844 & MATH\_ERREXCEPT) is non-zero, then the invalid floating-point exception |  
 3845 shall be raised. |

3846 **EXAMPLES**

3847 None.

3848 **APPLICATION USAGE**3849 On error, the expressions (math\_errhandling & MATH\_ERRNO) and (math\_errhandling &  
3850 MATH\_ERREXCEPT) are independent of each other, but at least one of them must be non-zero.3851 **RATIONALE**

3852 None.

3853 **FUTURE DIRECTIONS**

3854 None.

3855 **SEE ALSO**

3856 *cosh()*, *feclearexcept()*, *fetetestexcept()*, the Base Definitions volume of IEEE Std 1003.1-200x, Section |  
3857 4.18, Treatment of Error Conditions for Mathematical Functions, <math.h> |

3858 **CHANGE HISTORY**

3859 First released in Issue 4, Version 2.

3860 **Issue 5**

3861 Moved from X/OPEN UNIX extension to BASE.

3862 **Issue 6**3863 The *acosh()* function is no longer marked as an extension.

3864 The *acoshf()*, and *acoshl()* functions are added for alignment with the ISO/IEC 9899:1999  
3865 standard.

3866 The DESCRIPTION, RETURN VALUE, ERRORS, and APPLICATION USAGE sections are  
3867 revised to align with the ISO/IEC 9899:1999 standard.

3868 IEC 60559:1989 standard floating-point extensions over the ISO/IEC 9899:1999 standard are  
3869 marked.



3870 **NAME**

3871        *acosl* — arc cosine functions

3872 **SYNOPSIS**

3873        #include <math.h>

3874        long double *acosl*(long double *x*);

3875 **DESCRIPTION**

3876        Refer to *acos*().

3877 **NAME**3878 aio\_cancel — cancel an asynchronous I/O request (**REALTIME**)3879 **SYNOPSIS**

3880 AIO #include &lt;aio.h&gt;

3881 int aio\_cancel(int *fildev*, struct aiocb \**aiocbp*);

3882

3883 **DESCRIPTION**

3884 The *aio\_cancel()* function shall attempt to cancel one or more asynchronous I/O requests  
3885 currently outstanding against file descriptor *fildev*. The *aiocbp* argument points to the  
3886 asynchronous I/O control block for a particular request to be canceled. If *aiocbp* is NULL, then  
3887 all outstanding cancelable asynchronous I/O requests against *fildev* shall be canceled.

3888 Normal asynchronous notification shall occur for asynchronous I/O operations that are  
3889 successfully canceled. If there are requests that cannot be canceled, then the normal  
3890 asynchronous completion process shall take place for those requests when they are completed.

3891 For requested operations that are successfully canceled, the associated error status shall be set to  
3892 [ECANCELED] and the return status shall be -1. For requested operations that are not  
3893 successfully canceled, the *aiocbp* shall not be modified by *aio\_cancel()*.

3894 If *aiocbp* is not NULL, then if *fildev* does not have the same value as the file descriptor with which  
3895 the asynchronous operation was initiated, unspecified results occur.

3896 Which operations are cancelable is implementation-defined.

3897 **RETURN VALUE**

3898 The *aio\_cancel()* function shall return the value AIO\_CANCELED to the calling process if the  
3899 requested operation(s) were canceled. The value AIO\_NOTCANCELED shall be returned if at  
3900 least one of the requested operation(s) cannot be canceled because it is in progress. In this case,  
3901 the state of the other operations, if any, referenced in the call to *aio\_cancel()* is not indicated by  
3902 the return value of *aio\_cancel()*. The application may determine the state of affairs for these  
3903 operations by using *aio\_error()*. The value AIO\_ALLDONE is returned if all of the operations  
3904 have already completed. Otherwise, the function shall return -1 and set *errno* to indicate the  
3905 error.

3906 **ERRORS**

3907 The *aio\_cancel()* function shall fail if:

3908 [EBADF] The *fildev* argument is not a valid file descriptor.

3909 **EXAMPLES**

3910 None.

3911 **APPLICATION USAGE**

3912 The *aio\_cancel()* function is part of the Asynchronous Input and Output option and need not be  
3913 available on all implementations.

3914 **RATIONALE**

3915 None.

3916 **FUTURE DIRECTIONS**

3917 None.

3918 **SEE ALSO**

3919 *aio\_read()*, *aio\_write()*, the Base Definitions volume of IEEE Std 1003.1-200x, <**aio.h**>

3920 **CHANGE HISTORY**

3921 First released in Issue 5. Included for alignment with the POSIX Realtime Extension.

3922 **Issue 6**

3923 The [ENOSYS] error condition has been removed as stubs need not be provided if an  
3924 implementation does not support the Asynchronous Input and Output option.

3925 The APPLICATION USAGE section is added.

3926 **NAME**

3927 aio\_error — retrieve errors status for an asynchronous I/O operation (**REALTIME**)

3928 **SYNOPSIS**

3929 AIO `#include <aio.h>`

3930 `int aio_error(const struct aiocb *aiocbp);`

3931

3932 **DESCRIPTION**

3933 The *aio\_error()* function shall return the error status associated with the **aiocb** structure  
3934 referenced by the *aiocbp* argument. The error status for an asynchronous I/O operation is the  
3935 SIO *errno* value that would be set by the corresponding *read()*, *write()*, *fdatasync()*, or *fsync()*  
3936 operation. If the operation has not yet completed, then the error status shall be equal to  
3937 [EINPROGRESS].

3938 **RETURN VALUE**

3939 If the asynchronous I/O operation has completed successfully, then 0 shall be returned. If the  
3940 asynchronous operation has completed unsuccessfully, then the error status, as described for  
3941 SIO *read()*, *write()*, *fdatasync()*, and *fsync()*, shall be returned. If the asynchronous I/O operation has  
3942 not yet completed, then [EINPROGRESS] shall be returned.

3943 **ERRORS**

3944 The *aio\_error()* function may fail if:

3945 [EINVAL] The *aiocbp* argument does not refer to an asynchronous operation whose  
3946 return status has not yet been retrieved.

3947 **EXAMPLES**

3948 None.

3949 **APPLICATION USAGE**

3950 The *aio\_error()* function is part of the Asynchronous Input and Output option and need not be  
3951 available on all implementations.

3952 **RATIONALE**

3953 None.

3954 **FUTURE DIRECTIONS**

3955 None.

3956 **SEE ALSO**

3957 *aio\_cancel()*, *aio\_fsync()*, *aio\_read()*, *aio\_return()*, *aio\_write()*, *close()*, *exec*, *exit()*, *fork()*, *lio\_listio()*,  
3958 *lseek()*, *read()*, the Base Definitions volume of IEEE Std 1003.1-200x, **<aio.h>**

3959 **CHANGE HISTORY**

3960 First released in Issue 5. Included for alignment with the POSIX Realtime Extension.

3961 **Issue 6**

3962 The [ENOSYS] error condition has been removed as stubs need not be provided if an  
3963 implementation does not support the Asynchronous Input and Output option.

3964 The APPLICATION USAGE section is added.

3965 **NAME**3966 aio\_fsync — asynchronous file synchronization (**REALTIME**)3967 **SYNOPSIS**

3968 AIO #include &lt;aio.h&gt;

3969 int aio\_fsync(int op, struct aiocb \*aiocbp);

3970

3971 **DESCRIPTION**

3972 The *aio\_fsync()* function shall asynchronously force all I/O operations associated with the file |  
 3973 indicated by the file descriptor *aio\_fildes* member of the **aiocb** structure referenced by the *aiocbp* |  
 3974 argument and queued at the time of the call to *aio\_fsync()* to the synchronized I/O completion  
 3975 state. The function call shall return when the synchronization request has been initiated or  
 3976 queued to the file or device (even when the data cannot be synchronized immediately).

3977 If *op* is O\_DSYNC, all currently queued I/O operations shall be completed as if by a call to  
 3978 *fdatasync()*; that is, as defined for synchronized I/O data integrity completion. If *op* is O\_SYNC,  
 3979 all currently queued I/O operations shall be completed as if by a call to *fsync()*; that is, as  
 3980 defined for synchronized I/O file integrity completion. If the *aio\_fsync()* function fails, or if the  
 3981 operation queued by *aio\_fsync()* fails, then, as for *fsync()* and *fdatasync()*, outstanding I/O  
 3982 operations are not guaranteed to have been completed.

3983 If *aio\_fsync()* succeeds, then it is only the I/O that was queued at the time of the call to  
 3984 *aio\_fsync()* that is guaranteed to be forced to the relevant completion state. The completion of  
 3985 subsequent I/O on the file descriptor is not guaranteed to be completed in a synchronized  
 3986 fashion.

3987 The *aiocbp* argument refers to an asynchronous I/O control block. The *aiocbp* value may be used |  
 3988 as an argument to *aio\_error()* and *aio\_return()* in order to determine the error status and return  
 3989 status, respectively, of the asynchronous operation while it is proceeding. When the request is  
 3990 queued, the error status for the operation is [EINPROGRESS]. When all data has been  
 3991 successfully transferred, the error status shall be reset to reflect the success or failure of the  
 3992 operation. If the operation does not complete successfully, the error status for the operation shall  
 3993 be set to indicate the error. The *aio\_sigevent* member determines the asynchronous notification to  
 3994 occur as specified in Section 2.4.1 (on page 478) when all operations have achieved synchronized  
 3995 I/O completion. All other members of the structure referenced by *aiocbp* are ignored. If the  
 3996 control block referenced by *aiocbp* becomes an illegal address prior to asynchronous I/O  
 3997 completion, then the behavior is undefined.

3998 If the *aio\_fsync()* function fails or the *aiocbp* indicates an error condition, data is not guaranteed  
 3999 to have been successfully transferred.

4000 **RETURN VALUE**

4001 The *aio\_fsync()* function shall return the value 0 to the calling process if the I/O operation is  
 4002 successfully queued; otherwise, the function shall return the value -1 and set *errno* to indicate  
 4003 the error.

4004 **ERRORS**4005 The *aio\_fsync()* function shall fail if:

4006 [EAGAIN] The requested asynchronous operation was not queued due to temporary  
 4007 resource limitations.

4008 [EBADF] The *aio\_fildes* member of the **aiocb** structure referenced by the *aiocbp* argument  
 4009 is not a valid file descriptor open for writing.

- 4010 [EINVAL] This implementation does not support synchronized I/O for this file.
- 4011 [EINVAL] A value of *op* other than O\_DSYNC or O\_SYNC was specified.
- 4012 In the event that any of the queued I/O operations fail, *aio\_fsync()* shall return the error  
4013 condition defined for *read()* and *write()*. The error is returned in the error status for the  
4014 asynchronous *fsync()* operation, which can be retrieved using *aio\_error()*.
- 4015 **EXAMPLES**
- 4016 None.
- 4017 **APPLICATION USAGE**
- 4018 The *aio\_fsync()* function is part of the Asynchronous Input and Output option and need not be  
4019 available on all implementations.
- 4020 **RATIONALE**
- 4021 None.
- 4022 **FUTURE DIRECTIONS**
- 4023 None.
- 4024 **SEE ALSO**
- 4025 *fcntl()*, *fdatasync()*, *fsync()*, *open()*, *read()*, *write()*, the Base Definitions volume of  
4026 IEEE Std 1003.1-200x, <**aio.h**>
- 4027 **CHANGE HISTORY**
- 4028 First released in Issue 5. Included for alignment with the POSIX Realtime Extension.
- 4029 **Issue 6**
- 4030 The [ENOSYS] error condition has been removed as stubs need not be provided if an  
4031 implementation does not support the Asynchronous Input and Output option.
- 4032 The APPLICATION USAGE section is added.

4033 **NAME**4034 aio\_read — asynchronous read from a file (**REALTIME**)4035 **SYNOPSIS**

4036 AIO #include &lt;aio.h&gt;

4037 int aio\_read(struct aiocb \*aiocbp);

4038

4039 **DESCRIPTION**

4040 The *aio\_read()* function shall read *aiocbp->aio\_nbytes* from the file associated with *aiocbp->aio\_fildes* into the buffer pointed to by *aiocbp->aio\_buf*. The function call shall return when the read request has been initiated or queued to the file or device (even when the data cannot be delivered immediately).

4044 PIO If prioritized I/O is supported for this file, then the asynchronous operation shall be submitted at a priority equal to the scheduling priority of the process minus *aiocbp->aio\_reqprio*.

4046 The *aiocbp* value may be used as an argument to *aio\_error()* and *aio\_return()* in order to determine the error status and return status, respectively, of the asynchronous operation while it is proceeding. If an error condition is encountered during queuing, the function call shall return without having initiated or queued the request. The requested operation takes place at the absolute position in the file as given by *aio\_offset*, as if *lseek()* were called immediately prior to the operation with an *offset* equal to *aio\_offset* and a *whence* equal to *SEEK\_SET*. After a successful call to enqueue an asynchronous I/O operation, the value of the file offset for the file is unspecified.

4054 The *aiocbp->aio\_lio\_opcode* field shall be ignored by *aio\_read()*.

4055 The *aiocbp* argument points to an **aiocb** structure. If the buffer pointed to by *aiocbp->aio\_buf* or the control block pointed to by *aiocbp* becomes an illegal address prior to asynchronous I/O completion, then the behavior is undefined.

4058 Simultaneous asynchronous operations using the same *aiocbp* produce undefined results.

4059 SIO If synchronized I/O is enabled on the file associated with *aiocbp->aio\_fildes*, the behavior of this function shall be according to the definitions of synchronized I/O data integrity completion and synchronized I/O file integrity completion.

4062 For any system action that changes the process memory space while an asynchronous I/O is outstanding to the address range being changed, the result of that action is undefined.

4064 For regular files, no data transfer shall occur past the offset maximum established in the open file description associated with *aiocbp->aio\_fildes*.

4066 **RETURN VALUE**

4067 The *aio\_read()* function shall return the value zero to the calling process if the I/O operation is successfully queued; otherwise, the function shall return the value  $-1$  and set *errno* to indicate the error.

4070 **ERRORS**

4071 The *aio\_read()* function shall fail if:

4072 [EAGAIN] The requested asynchronous I/O operation was not queued due to system resource limitations.

4074 Each of the following conditions may be detected synchronously at the time of the call to *aio\_read()*, or asynchronously. If any of the conditions below are detected synchronously, the *aio\_read()* function shall return  $-1$  and set *errno* to the corresponding value. If any of the conditions below are detected asynchronously, the return status of the asynchronous operation

- 4078 is set to `-1`, and the error status of the asynchronous operation is set to the corresponding value.
- 4079 [EBADF] The `aiocbp->aio_fildes` argument is not a valid file descriptor open for reading.
- 4080 [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.
- 4081
- 4082 In the case that the `aio_read()` successfully queues the I/O operation but the operation is
- 4083 subsequently canceled or encounters an error, the return status of the asynchronous operation is
- 4084 one of the values normally returned by the `read()` function call. In addition, the error status of
- 4085 the asynchronous operation is set to one of the error statuses normally set by the `read()` function
- 4086 call, or one of the following values:
- 4087 [EBADF] The `aiocbp->aio_fildes` argument is not a valid file descriptor open for reading.
- 4088 [ECANCELED] The requested I/O was canceled before the I/O completed due to an explicit
- 4089 `aio_cancel()` request.
- 4090 [EINVAL] The file offset value implied by `aiocbp->aio_offset` would be invalid.
- 4091 The following condition may be detected synchronously or asynchronously:
- 4092 [EOVERFLOW] The file is a regular file, `aiocbp->aio_nbytes` is greater than 0, and the starting
- 4093 offset in `aiocbp->aio_offset` is before the end-of-file and is at or beyond the offset
- 4094 maximum in the open file description associated with `aiocbp->aio_fildes`.

#### 4095 EXAMPLES

4096 None.

#### 4097 APPLICATION USAGE

4098 The `aio_read()` function is part of the Asynchronous Input and Output option and need not be

4099 available on all implementations.

#### 4100 RATIONALE

4101 None.

#### 4102 FUTURE DIRECTIONS

4103 None.

#### 4104 SEE ALSO

4105 `aio_cancel()`, `aio_error()`, `lio_listio()`, `aio_return()`, `aio_write()`, `close()`, `exec`, `exit()`, `fork()`, `lseek()`,

4106 `read()`, the Base Definitions volume of IEEE Std 1003.1-200x, <**aio.h**>

#### 4107 CHANGE HISTORY

4108 First released in Issue 5. Included for alignment with the POSIX Realtime Extension.

4109 Large File Summit extensions are added.

#### 4110 Issue 6

4111 The [ENOSYS] error condition has been removed as stubs need not be provided if an

4112 implementation does not support the Asynchronous Input and Output option.

4113 The APPLICATION USAGE section is added. |



- 4114 The following new requirements on POSIX implementations derive from alignment with the |  
4115 Single UNIX Specification:
- 4116 • In the DESCRIPTION, text is added to indicate setting of the offset maximum in the open file  
4117 description. This change is to support large files.
  - 4118 • In the ERRORS section, the [Eoverflow] condition is added. This change is to support  
4119 large files.

4120 **NAME**4121 aio\_return — retrieve return status of an asynchronous I/O operation (**REALTIME**)4122 **SYNOPSIS**4123 AIO `#include <aio.h>`4124 `ssize_t aio_return(struct aiocb *aiocbp);`

4125

4126 **DESCRIPTION**

4127 The `aio_return()` function shall return the return status associated with the **aiocb** structure  
4128 referenced by the `aiocbp` argument. The return status for an asynchronous I/O operation is the  
4129 value that would be returned by the corresponding `read()`, `write()`, or `fsync()` function call. If the  
4130 error status for the operation is equal to [EINPROGRESS], then the return status for the  
4131 operation is undefined. The `aio_return()` function may be called exactly once to retrieve the  
4132 return status of a given asynchronous operation; thereafter, if the same **aiocb** structure is used in  
4133 a call to `aio_return()` or `aio_error()`, an error may be returned. When the **aiocb** structure referred  
4134 to by `aiocbp` is used to submit another asynchronous operation, then `aio_return()` may be  
4135 successfully used to retrieve the return status of that operation.

4136 **RETURN VALUE**

4137 If the asynchronous I/O operation has completed, then the return status, as described for `read()`,  
4138 `write()`, and `fsync()`, shall be returned. If the asynchronous I/O operation has not yet completed,  
4139 the results of `aio_return()` are undefined.

4140 **ERRORS**4141 The `aio_return()` function may fail if:

4142 [EINVAL] The `aiocbp` argument does not refer to an asynchronous operation whose  
4143 return status has not yet been retrieved.

4144 **EXAMPLES**

4145 None.

4146 **APPLICATION USAGE**

4147 The `aio_return()` function is part of the Asynchronous Input and Output option and need not be  
4148 available on all implementations.

4149 **RATIONALE**

4150 None.

4151 **FUTURE DIRECTIONS**

4152 None.

4153 **SEE ALSO**

4154 `aio_cancel()`, `aio_error()`, `aio_fsync()`, `aio_read()`, `aio_write()`, `close()`, `exec`, `exit()`, `fork()`, `lio_listio()`,  
4155 `lseek()`, `read()`, the Base Definitions volume of IEEE Std 1003.1-200x, **<aio.h>**

4156 **CHANGE HISTORY**

4157 First released in Issue 5. Included for alignment with the POSIX Realtime Extension.

4158 **Issue 6**

4159 The [ENOSYS] error condition has been removed as stubs need not be provided if an  
4160 implementation does not support the Asynchronous Input and Output option.

4161 The APPLICATION USAGE section is added.

4162 The [EINVAL] error condition is updated as a “may fail”. This is for consistency with the  
4163 DESCRIPTION.

4164 **NAME**4165 aio\_suspend — wait for an asynchronous I/O request (**REALTIME**)4166 **SYNOPSIS**

```
4167 AIO #include <aio.h>
4168 int aio_suspend(const struct aiocb * const list[], int nent,
4169               const struct timespec *timeout);
4170
```

4171 **DESCRIPTION**

4172 The *aio\_suspend()* function shall suspend the calling thread until at least one of the asynchronous  
 4173 I/O operations referenced by the *list* argument has completed, until a signal interrupts the  
 4174 function, or, if *timeout* is not NULL, until the time interval specified by *timeout* has passed. If any  
 4175 of the **aiocb** structures in the list correspond to completed asynchronous I/O operations (that is,  
 4176 the error status for the operation is not equal to [EINPROGRESS]) at the time of the call, the  
 4177 function shall return without suspending the calling thread. The *list* argument is an array of  
 4178 pointers to asynchronous I/O control blocks. The *nent* argument indicates the number of  
 4179 elements in the array. Each **aiocb** structure pointed to has been used in initiating an  
 4180 asynchronous I/O request via *aio\_read()*, *aio\_write()*, or *lio\_listio()*. This array may contain  
 4181 NULL pointers, which are ignored. If this array contains pointers that refer to **aiocb** structures  
 4182 that have not been used in submitting asynchronous I/O, the effect is undefined.

4183 If the time interval indicated in the **timespec** structure pointed to by *timeout* passes before any of  
 4184 the I/O operations referenced by *list* are completed, then *aio\_suspend()* shall return with an  
 4185 error. If the Monotonic Clock option is supported, the clock that shall be used to measure this  
 4186 time interval shall be the CLOCK\_MONOTONIC clock.

4187 **RETURN VALUE**

4188 If the *aio\_suspend()* function returns after one or more asynchronous I/O operations have  
 4189 completed, the function shall return zero. Otherwise, the function shall return a value of -1 and  
 4190 set *errno* to indicate the error.

4191 The application may determine which asynchronous I/O completed by scanning the associated  
 4192 error and return status using *aio\_error()* and *aio\_return()*, respectively.

4193 **ERRORS**

4194 The *aio\_suspend()* function shall fail if:

4195 [EAGAIN] No asynchronous I/O indicated in the list referenced by *list* completed in the  
 4196 time interval indicated by *timeout*.

4197 [EINTR] A signal interrupted the *aio\_suspend()* function. Note that, since each  
 4198 asynchronous I/O operation may possibly provoke a signal when it  
 4199 completes, this error return may be caused by the completion of one (or more)  
 4200 of the very I/O operations being awaited.

4201 **EXAMPLES**

4202 None.

4203 **APPLICATION USAGE**

4204 The *aio\_suspend()* function is part of the Asynchronous Input and Output option and need not  
 4205 be available on all implementations.

4206 **RATIONALE**

4207 None.

4208 **FUTURE DIRECTIONS**

4209       None.

4210 **SEE ALSO**

4211       *aio\_read()*, *aio\_write()*, *lio\_listio()*, the Base Definitions volume of IEEE Std 1003.1-200x, <**aio.h**>

4212 **CHANGE HISTORY**

4213       First released in Issue 5. Included for alignment with the POSIX Realtime Extension.

4214 **Issue 6**

4215       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.

4217       The APPLICATION USAGE section is added.

4218       The DESCRIPTION is updated for alignment with IEEE Std 1003.1j-2000 by specifying that the CLOCK\_MONOTONIC clock, if supported, is used.

4220 **NAME**4221 aio\_write — asynchronous write to a file (**REALTIME**)4222 **SYNOPSIS**

4223 AIO #include &lt;aio.h&gt;

4224 int aio\_write(struct aiocb \*aiocbp);

4225

4226 **DESCRIPTION**

4227 The *aio\_write()* function shall write *aiocbp->aio\_nbytes* to the file associated with *aiocbp->aio\_fildes* |  
 4228 from the buffer pointed to by *aiocbp->aio\_buf*. The function shall return when the write request |  
 4229 has been initiated or, at a minimum, queued to the file or device. |

4230 PIO If prioritized I/O is supported for this file, then the asynchronous operation shall be submitted |  
 4231 at a priority equal to the scheduling priority of the process minus *aiocbp->aio\_reqprio*. |

4232 The *aiocbp* may be used as an argument to *aio\_error()* and *aio\_return()* in order to determine the |  
 4233 error status and return status, respectively, of the asynchronous operation while it is proceeding.

4234 The *aiocbp* argument points to an **aiocb** structure. If the buffer pointed to by *aiocbp->aio\_buf* or |  
 4235 the control block pointed to by *aiocbp* becomes an illegal address prior to asynchronous I/O |  
 4236 completion, then the behavior is undefined.

4237 If **O\_APPEND** is not set for the file descriptor *aio\_fildes*, then the requested operation shall take |  
 4238 place at the absolute position in the file as given by *aio\_offset*, as if *lseek()* were called |  
 4239 immediately prior to the operation with an *offset* equal to *aio\_offset* and a *whence* equal to |  
 4240 **SEEK\_SET**. If **O\_APPEND** is set for the file descriptor, write operations append to the file in the |  
 4241 same order as the calls were made. After a successful call to enqueue an asynchronous I/O |  
 4242 operation, the value of the file offset for the file is unspecified.

4243 The *aiocbp->aio\_lio\_opcode* field shall be ignored by *aio\_write()*.

4244 Simultaneous asynchronous operations using the same *aiocbp* produce undefined results.

4245 SIO If synchronized I/O is enabled on the file associated with *aiocbp->aio\_fildes*, the behavior of this |  
 4246 function shall be according to the definitions of synchronized I/O data integrity completion, and |  
 4247 synchronized I/O file integrity completion.

4248 For any system action that changes the process memory space while an asynchronous I/O is |  
 4249 outstanding to the address range being changed, the result of that action is undefined.

4250 For regular files, no data transfer shall occur past the offset maximum established in the open |  
 4251 file description associated with *aiocbp->aio\_fildes*.

4252 **RETURN VALUE**

4253 The *aio\_write()* function shall return the value zero to the calling process if the I/O operation is |  
 4254 successfully queued; otherwise, the function shall return the value  $-1$  and set *errno* to indicate |  
 4255 the error.

4256 **ERRORS**

4257 The *aio\_write()* function shall fail if:

4258 [EAGAIN] The requested asynchronous I/O operation was not queued due to system |  
 4259 resource limitations.

4260 Each of the following conditions may be detected synchronously at the time of the call to |  
 4261 *aio\_write()*, or asynchronously. If any of the conditions below are detected synchronously, the |  
 4262 *aio\_write()* function shall return  $-1$  and set *errno* to the corresponding value. If any of the |  
 4263 conditions below are detected asynchronously, the return status of the asynchronous operation

4264 shall be set to -1, and the error status of the asynchronous operation is set to the corresponding  
4265 value.

4266 [EBADF] The *aiocbp->aio\_fildes* argument is not a valid file descriptor open for writing.

4267 [EINVAL] The file offset value implied by *aiocbp->aio\_offset* would be invalid, *aiocbp->aio\_reqprio*  
4268 is not a valid value, or *aiocbp->aio\_nbytes* is an invalid value.

4269 In the case that the *aio\_write()* successfully queues the I/O operation, the return status of the  
4270 asynchronous operation shall be one of the values normally returned by the *write()* function call.  
4271 If the operation is successfully queued but is subsequently canceled or encounters an error, the  
4272 error status for the asynchronous operation contains one of the values normally set by the  
4273 *write()* function call, or one of the following:

4274 [EBADF] The *aiocbp->aio\_fildes* argument is not a valid file descriptor open for writing.

4275 [EINVAL] The file offset value implied by *aiocbp->aio\_offset* would be invalid.

4276 [ECANCELED] The requested I/O was canceled before the I/O completed due to an explicit  
4277 *aio\_cancel()* request.

4278 The following condition may be detected synchronously or asynchronously:

4279 [EFBIG] The file is a regular file, *aiocbp->aio\_nbytes* is greater than 0, and the starting  
4280 offset in *aiocbp->aio\_offset* is at or beyond the offset maximum in the open file  
4281 description associated with *aiocbp->aio\_fildes*.

#### 4282 EXAMPLES

4283 None.

#### 4284 APPLICATION USAGE

4285 The *aio\_write()* function is part of the Asynchronous Input and Output option and need not be  
4286 available on all implementations.

#### 4287 RATIONALE

4288 None.

#### 4289 FUTURE DIRECTIONS

4290 None.

#### 4291 SEE ALSO

4292 *aio\_cancel()*, *aio\_error()*, *aio\_read()*, *aio\_return()*, *close()*, *exec*, *exit()*, *fork()*, *lio\_listio()*, *lseek()*,  
4293 *write()*, the Base Definitions volume of IEEE Std 1003.1-200x, <**aio.h**>

#### 4294 CHANGE HISTORY

4295 First released in Issue 5. Included for alignment with the POSIX Realtime Extension.

4296 Large File Summit extensions are added.

#### 4297 Issue 6

4298 The [ENOSYS] error condition has been removed as stubs need not be provided if an  
4299 implementation does not support the Asynchronous Input and Output option.

4300 The APPLICATION USAGE section is added.

4301 The following new requirements on POSIX implementations derive from alignment with the  
4302 Single UNIX Specification:

- 4303 • In the DESCRIPTION, text is added to indicate that for regular files no data transfer occurs  
4304 past the offset maximum established in the open file description associated with *aiocbp->aio\_fildes*.

4305

4306

- The [EFBIG] error is added as part of the large file support extensions.

4307 **NAME**

4308 alarm — schedule an alarm signal

4309 **SYNOPSIS**

4310 #include &lt;unistd.h&gt;

4311 unsigned alarm(unsigned *seconds*);4312 **DESCRIPTION**

4313 The *alarm()* function shall cause the system to generate a SIGALRM signal for the process after  
4314 the number of realtime seconds specified by *seconds* have elapsed. Processor scheduling delays  
4315 may prevent the process from handling the signal as soon as it is generated.

4316 If *seconds* is 0, a pending alarm request, if any, is canceled.

4317 Alarm requests are not stacked; only one SIGALRM generation can be scheduled in this manner.  
4318 If the SIGALRM signal has not yet been generated, the call shall result in rescheduling the time  
4319 at which the SIGALRM signal is generated.

4320 XSI Interactions between *alarm()* and any of *setitimer()*, *ualarm()*, or *usleep()* are unspecified.

4321 **RETURN VALUE**

4322 If there is a previous *alarm()* request with time remaining, *alarm()* shall return a non-zero value  
4323 that is the number of seconds until the previous request would have generated a SIGALRM  
4324 signal. Otherwise, *alarm()* shall return 0.

4325 **ERRORS**

4326 The *alarm()* function is always successful, and no return value is reserved to indicate an error.

4327 **EXAMPLES**

4328 None.

4329 **APPLICATION USAGE**

4330 The *fork()* function clears pending alarms in the child process. A new process image created by  
4331 one of the *exec* functions inherits the time left to an alarm signal in the old process' image.

4332 Application writers should note that the type of the argument *seconds* and the return value of  
4333 *alarm()* is **unsigned**. That means that a Strictly Conforming POSIX System Interfaces  
4334 Application cannot pass a value greater than the minimum guaranteed value for {UINT\_MAX},  
4335 which the ISO C standard sets as 65 535, and any application passing a larger value is restricting  
4336 its portability. A different type was considered, but historical implementations, including those  
4337 with a 16-bit **int** type, consistently use either **unsigned** or **int**.

4338 Application writers should be aware of possible interactions when the same process uses both  
4339 the *alarm()* and *sleep()* functions.

4340 **RATIONALE**

4341 Many historical implementations (including Version 7 and System V) allow an alarm to occur up  
4342 to a second early. Other implementations allow alarms up to half a second or one clock tick  
4343 early or do not allow them to occur early at all. The latter is considered most appropriate, since it  
4344 gives the most predictable behavior, especially since the signal can always be delayed for an  
4345 indefinite amount of time due to scheduling. Applications can thus choose the *seconds* argument  
4346 as the minimum amount of time they wish to have elapse before the signal.

4347 The term *realtime* here and elsewhere (*sleep()*, *times()*) is intended to mean “wall clock” time as  
4348 common English usage, and has nothing to do with “realtime operating systems”. It is in  
4349 contrast to *virtual time*, which could be misinterpreted if just *time* were used.

4350 In some implementations, including 4.3 BSD, very large values of the *seconds* argument are  
4351 silently rounded down to an implementation-defined maximum value. This maximum is large



4352 enough (on the order of several months) that the effect is not noticeable.

4353 There were two possible choices for alarm generation in multi-threaded applications: generation  
4354 for the calling thread or generation for the process. The first option would not have been  
4355 particularly useful since the alarm state is maintained on a per-process basis and the alarm that  
4356 is established by the last invocation of *alarm()* is the only one that would be active.

4357 Furthermore, allowing generation of an asynchronous signal for a thread would have introduced  
4358 an exception to the overall signal model. This requires a compelling reason in order to be  
4359 justified.

#### 4360 **FUTURE DIRECTIONS**

4361 None.

#### 4362 **SEE ALSO**

4363 *alarm()*, *exec*, *fork()*, *getitimer()*, *pause()*, *sigaction()*, *sleep()*, *ualarm()*, *usleep()*, the Base  
4364 Definitions volume of IEEE Std 1003.1-200x, <signal.h>, <unistd.h>

#### 4365 **CHANGE HISTORY**

4366 First released in Issue 1. Derived from Issue 1 of the SVID.

#### 4367 **Issue 6**

4368 The following new requirements on POSIX implementations derive from alignment with the  
4369 Single UNIX Specification:

- 4370 • The DESCRIPTION is updated to indicate that interactions with the *setitimer()*, *ualarm()*, and  
4371 *usleep()* functions are unspecified.

## 4372 NAME

4373 asctime, asctime\_r — convert date and time to a string

## 4374 SYNOPSIS

4375 #include &lt;time.h&gt;

4376 char \*asctime(const struct tm \*timeptr);

4377 TSF char \*asctime\_r(const struct tm \*restrict tm, char \*restrict buf);

4378

## 4379 DESCRIPTION

4380 CX For *asctime()*: The functionality described on this reference page is aligned with the ISO C |  
 4381 standard. Any conflict between the requirements described here and the ISO C standard is  
 4382 unintentional. This volume of IEEE Std 1003.1-200x defers to the ISO C standard.

4383 The *asctime()* function shall convert the broken-down time in the structure pointed to by *timeptr*  
 4384 into a string in the form:

4385 Sun Sep 16 01:03:52 1973\n\0

4386 using the equivalent of the following algorithm:

```

4387 char *asctime(const struct tm *timeptr)
4388 {
4389     static char wday_name[7][3] = {
4390         "Sun", "Mon", "Tue", "Wed", "Thu", "Fri", "Sat"
4391     };
4392     static char mon_name[12][3] = {
4393         "Jan", "Feb", "Mar", "Apr", "May", "Jun",
4394         "Jul", "Aug", "Sep", "Oct", "Nov", "Dec"
4395     };
4396     static char result[26];
4397
4398     sprintf(result, "%.3s %.3s%3d %.2d:%.2d:%.2d %d\n",
4399             wday_name[timeptr->tm_wday],
4400             mon_name[timeptr->tm_mon],
4401             timeptr->tm_mday, timeptr->tm_hour,
4402             timeptr->tm_min, timeptr->tm_sec,
4403             1900 + timeptr->tm_year);
4404     return result;
4405 }
```

4405 The **tm** structure is defined in the <time.h> header. |

4406 CX The *asctime()*, *ctime()*, *gmtime()*, and *localtime()* functions shall return values in one of two static  
 4407 objects: a broken-down time structure and an array of type **char**. Execution of any of the  
 4408 functions may overwrite the information returned in either of these objects by any of the other  
 4409 functions.

4410 The *asctime()* function need not be reentrant. A function that is not required to be reentrant is not  
 4411 required to be thread-safe.

4412 TSF The *asctime\_r()* function shall convert the broken-down time in the structure pointed to by *tm*  
 4413 into a string (of the same form as that returned by *asctime()*) that is placed in the user-supplied  
 4414 buffer pointed to by *buf* (which shall contain at least 26 bytes) and then return *buf*.

4415 **RETURN VALUE**

4416 Upon successful completion, *asctime()* shall return a pointer to the string.

4417 TSF Upon successful completion, *asctime\_r()* shall return a pointer to a character string containing  
4418 the date and time. This string is pointed to by the argument *buf*. If the function is unsuccessful,  
4419 it shall return NULL.

4420 **ERRORS**

4421 No errors are defined.

4422 **EXAMPLES**

4423 None.

4424 **APPLICATION USAGE**

4425 Values for the broken-down time structure can be obtained by calling *gmtime()* or *localtime()*.  
4426 This function is included for compatibility with older implementations, and does not support  
4427 localized date and time formats. Applications should use *strptime()* to achieve maximum  
4428 portability.

4429 The *asctime\_r()* function is thread-safe and shall return values in a user-supplied buffer instead  
4430 of possibly using a static data area that may be overwritten by each call.

4431 **RATIONALE**

4432 None.

4433 **FUTURE DIRECTIONS**

4434 None.

4435 **SEE ALSO**

4436 *clock()*, *ctime()*, *difftime()*, *gmtime()*, *localtime()*, *mktime()*, *strptime()*, *strptime()*, *time()*, *utime()*,  
4437 the Base Definitions volume of IEEE Std 1003.1-200x, <**time.h**>

4438 **CHANGE HISTORY**

4439 First released in Issue 1. Derived from Issue 1 of the SVID.

4440 **Issue 5**

4441 Normative text previously in the APPLICATION USAGE section is moved to the  
4442 DESCRIPTION.

4443 The *asctime\_r()* function is included for alignment with the POSIX Threads Extension.

4444 A note indicating that the *asctime()* function need not be reentrant is added to the  
4445 DESCRIPTION.

4446 **Issue 6**

4447 The *asctime\_r()* function is marked as part of the Thread-Safe Functions option.

4448 Extensions beyond the ISO C standard are now marked.

4449 The APPLICATION USAGE section is updated to include a note on the thread-safe function and  
4450 its avoidance of possibly using a static data area.

4451 The DESCRIPTION of *asctime\_r()* is updated to describe the format of the string returned.

4452 The **restrict** keyword is added to the *asctime\_r()* prototype for alignment with the  
4453 ISO/IEC 9899:1999 standard.

4454 **NAME**

4455 asin, asinf, asinl — arc sine function

4456 **SYNOPSIS**

4457 #include &lt;math.h&gt;

4458 double asin(double x);

4459 float asinf(float x);

4460 long double asinl(long double x);

4461 **DESCRIPTION**

4462 **CX** The functionality described on this reference page is aligned with the ISO C standard. Any  
 4463 conflict between the requirements described here and the ISO C standard is unintentional. This  
 4464 volume of IEEE Std 1003.1-200x defers to the ISO C standard.

4465 These functions shall compute the principal value of the arc sine of their argument *x*. The value  
 4466 of *x* should be in the range  $[-1,1]$ .

4467 An application wishing to check for error situations should set *errno* to zero and call  
 4468 *feclearexcept*(FE\_ALL\_EXCEPT) before calling these functions. On return, if *errno* is non-zero or  
 4469 *fetetestexcept*(FE\_INVALID | FE\_DIVBYZERO | FE\_OVERFLOW | FE\_UNDERFLOW) is non-  
 4470 zero, an error has occurred.

4471 **RETURN VALUE**

4472 Upon successful completion, these functions shall return the arc sine of *x*, in the range  
 4473  $[-\pi/2, \pi/2]$  radians.

4474 **MX** For finite values of *x* not in the range  $[-1,1]$ , a domain error shall occur, and either a NaN (if  
 4475 supported), or an implementation-defined value shall be returned.

4476 **MX** If *x* is NaN, a NaN shall be returned.

4477 If *x* is  $\pm 0$ , *x* shall be returned.

4478 If *x* is  $\pm \text{Inf}$ , a domain error shall occur, and either a NaN (if supported), or an implementation-  
 4479 defined value shall be returned.

4480 If *x* is subnormal, a range error may occur and *x* should be returned.

4481 **ERRORS**

4482 These functions shall fail if:

4483 **MX** **Domain Error** The *x* argument is finite and is not in the range  $[-1,1]$ , or is  $\pm \text{Inf}$ .

4484 If the integer expression (math\_errhandling & MATH\_ERRNO) is non-zero, |  
 4485 then *errno* shall be set to [EDOM]. If the integer expression (math\_errhandling |  
 4486 & MATH\_ERREXCEPT) is non-zero, then the invalid floating-point exception |  
 4487 shall be raised. |

4488 These functions may fail if:

4489 **MX** **Range Error** The value of *x* is subnormal.

4490 If the integer expression (math\_errhandling & MATH\_ERRNO) is non-zero, |  
 4491 then *errno* shall be set to [ERANGE]. If the integer expression |  
 4492 (math\_errhandling & MATH\_ERREXCEPT) is non-zero, then the underflow |  
 4493 floating-point exception shall be raised. |

4494 **EXAMPLES**

4495       None.

4496 **APPLICATION USAGE**

4497       On error, the expressions (`math_errhandling & MATH_ERRNO`) and (`math_errhandling &`  
4498       `MATH_ERREXCEPT`) are independent of each other, but at least one of them must be non-zero.

4499 **RATIONALE**

4500       None.

4501 **FUTURE DIRECTIONS**

4502       None.

4503 **SEE ALSO**

4504       *feclearexcept()*, *fetestexcept()*, *isnan()*, *sin()*, the Base Definitions volume of IEEE Std 1003.1-200x, |  
4505       Section 4.18, Treatment of Error Conditions for Mathematical Functions, <**math.h**> |

4506 **CHANGE HISTORY**

4507       First released in Issue 1. Derived from Issue 1 of the SVID.

4508 **Issue 5**

4509       The DESCRIPTION is updated to indicate how an application should check for an error. This  
4510       text was previously published in the APPLICATION USAGE section.

4511 **Issue 6**4512       The *asinf()* and *asinl()* functions are added for alignment with the ISO/IEC 9899:1999 standard.

4513       The DESCRIPTION, RETURN VALUE, ERRORS, and APPLICATION USAGE sections are  
4514       revised to align with the ISO/IEC 9899:1999 standard.

4515       IEC 60559:1989 standard floating-point extensions over the ISO/IEC 9899:1999 standard are  
4516       marked.

4517 **NAME**4518        **asinf** — arc sine function4519 **SYNOPSIS**

4520        #include &lt;math.h&gt;

4521        float asinf(float x);

4522 **DESCRIPTION**4523        Refer to *asin()*.

4524 **NAME**

4525        asinhf, asinhl — inverse hyperbolic sine functions

4526 **SYNOPSIS**

4527        #include &lt;math.h&gt;

4528        float asinhf(float x);

4529        long double asinhl(long double x);

4530 **DESCRIPTION**4531        Refer to *asinh()*.

4532 **NAME**

4533 asinh, asinhf, asinhl — inverse hyperbolic sine functions

4534 **SYNOPSIS**

```
4535 #include <math.h>
4536 double asinh(double x);
4537 float asinhf(float x);
4538 long double asinhl(long double x);
```

4539 **DESCRIPTION**

4540 **CX** The functionality described on this reference page is aligned with the ISO C standard. Any  
 4541 conflict between the requirements described here and the ISO C standard is unintentional. This  
 4542 volume of IEEE Std 1003.1-200x defers to the ISO C standard.

4543 These functions shall compute the inverse hyperbolic sine of their argument *x*.

4544 An application wishing to check for error situations should set *errno* to zero and call  
 4545 *feclearexcept*(FE\_ALL\_EXCEPT) before calling these functions. On return, if *errno* is non-zero or  
 4546 *fetetestexcept*(FE\_INVALID | FE\_DIVBYZERO | FE\_OVERFLOW | FE\_UNDERFLOW) is non-  
 4547 zero, an error has occurred.

4548 **RETURN VALUE**

4549 Upon successful completion, these functions shall return the inverse hyperbolic sine, of their  
 4550 argument.

4551 **MX** If *x* is NaN, a NaN shall be returned.

4552 If *x* is  $\pm 0$ , or  $\pm \text{Inf}$ , *x* shall be returned.

4553 If *x* is subnormal, a range error may occur and *x* should be returned.

4554 **ERRORS**

4555 These functions may fail if:

4556 **MX** **Range Error** The value of *x* is subnormal.

4557 If the integer expression (math\_errhandling & MATH\_ERRNO) is non-zero, |  
 4558 then *errno* shall be set to [ERANGE]. If the integer expression |  
 4559 (math\_errhandling & MATH\_ERREXCEPT) is non-zero, then the underflow |  
 4560 floating-point exception shall be raised. |

4561 **EXAMPLES**

4562 None.

4563 **APPLICATION USAGE**

4564 On error, the expressions (math\_errhandling & MATH\_ERRNO) and (math\_errhandling &  
 4565 MATH\_ERREXCEPT) are independent of each other, but at least one of them must be non-zero.

4566 **RATIONALE**

4567 None.

4568 **FUTURE DIRECTIONS**

4569 None.

4570 **SEE ALSO**

4571 *feclearexcept*(), *fetetestexcept*(), *sinh*(), the Base Definitions volume of IEEE Std 1003.1-200x, Section |  
 4572 4.18, Treatment of Error Conditions for Mathematical Functions, <math.h> |



4573 **CHANGE HISTORY**

4574 First released in Issue 4, Version 2.

4575 **Issue 5**

4576 Moved from X/OPEN UNIX extension to BASE.

4577 **Issue 6**4578 The *asinh()* function is no longer marked as an extension.4579 The *asinhf()*, and *asinhf\_l()* functions are added for alignment with the ISO/IEC 9899:1999  
4580 standard.4581 The DESCRIPTION, RETURN VALUE, ERRORS, and APPLICATION USAGE sections are  
4582 revised to align with the ISO/IEC 9899:1999 standard.4583 IEC 60559:1989 standard floating-point extensions over the ISO/IEC 9899:1999 standard are  
4584 marked.

4585 **NAME**4586        **asinl** — arc sine function4587 **SYNOPSIS**

4588        #include &lt;math.h&gt;

4589        long double asinl(long double x);

4590 **DESCRIPTION**4591        Refer to *asin()*.

4592 **NAME**

4593        assert — insert program diagnostics

4594 **SYNOPSIS**

4595        #include &lt;assert.h&gt;

4596        void assert(*scalar expression*);4597 **DESCRIPTION**

4598 **cx**     The functionality described on this reference page is aligned with the ISO C standard. Any  
4599 conflict between the requirements described here and the ISO C standard is unintentional. This  
4600 volume of IEEE Std 1003.1-200x defers to the ISO C standard.

4601     The *assert()* macro shall insert diagnostics into programs; it shall expand to a **void** expression. |  
4602     When it is executed, if *expression* (which shall have a **scalar** type) is false (that is, compares equal  
4603     to 0), *assert()* shall write information about the particular call that failed on *stderr* and shall call  
4604     *abort()*.

4605     The information written about the call that failed shall include the text of the argument, the  
4606     name of the source file, the source file line number, and the name of the enclosing function, the  
4607     latter are, respectively, the values of the preprocessing macros `__FILE__` and `__LINE__` and of  
4608     the identifier `__func__`.

4609     Forcing a definition of the name `NDEBUG`, either from the compiler command line or with the  
4610     preprocessor control statement `#define NDEBUG` ahead of the `#include <assert.h>` statement,  
4611     shall stop assertions from being compiled into the program.

4612 **RETURN VALUE**4613        The *assert()* macro shall not return a value.4614 **ERRORS**

4615        No errors are defined.

4616 **EXAMPLES**

4617        None.

4618 **APPLICATION USAGE**

4619        None.

4620 **RATIONALE**

4621        None.

4622 **FUTURE DIRECTIONS**

4623        None.

4624 **SEE ALSO**4625        *abort()*, the Base Definitions volume of IEEE Std 1003.1-200x, `<assert.h>`, *stderr*4626 **CHANGE HISTORY**

4627        First released in Issue 1. Derived from Issue 1 of the SVID.

4628 **Issue 6**

4629     The prototype for the *expression* argument to *assert()* is changed from **int** to **scalar** for alignment  
4630     with the ISO/IEC 9899:1999 standard.

4631     The DESCRIPTION of *assert()* is updated for alignment with the ISO/IEC 9899:1999 standard.

4632 **NAME**

4633 atan, atanf, atanl — arc tangent function

4634 **SYNOPSIS**

4635 #include &lt;math.h&gt;

4636 double atan(double x);

4637 float atanf(float x);

4638 long double atanl(long double x);

4639 **DESCRIPTION**

4640 **CX** The functionality described on this reference page is aligned with the ISO C standard. Any  
 4641 conflict between the requirements described here and the ISO C standard is unintentional. This  
 4642 volume of IEEE Std 1003.1-200x defers to the ISO C standard.

4643 These functions shall compute the principal value of the arc tangent of their argument  $x$ .

4644 An application wishing to check for error situations should set *errno* to zero and call  
 4645 *feclearexcept*(FE\_ALL\_EXCEPT) before calling these functions. On return, if *errno* is non-zero or  
 4646 *fetestexcept*(FE\_INVALID | FE\_DIVBYZERO | FE\_OVERFLOW | FE\_UNDERFLOW) is non-  
 4647 zero, an error has occurred.

4648 **RETURN VALUE**4649 Upon successful completion, these functions shall return the arc tangent of  $x$  in the range  
4650  $[-\pi/2, \pi/2]$  radians.4651 **MX** If  $x$  is NaN, a NaN shall be returned.4652 If  $x$  is  $\pm 0$   $x$  shall be returned.4653 If  $x$  is  $\pm\text{Inf}$ ,  $\pm\pi/2$  shall be returned.4654 If  $x$  is subnormal, a range error may occur and  $x$  should be returned.4655 **ERRORS**

4656 These functions may fail if:

4657 **MX** **Range Error** The value of  $x$  is subnormal.

4658 If the integer expression (math\_errhandling & MATH\_ERRNO) is non-zero, |  
 4659 then *errno* shall be set to [ERANGE]. If the integer expression |  
 4660 (math\_errhandling & MATH\_ERREXCEPT) is non-zero, then the underflow |  
 4661 floating-point exception shall be raised. |

4662 **EXAMPLES**

4663 None.

4664 **APPLICATION USAGE**4665 On error, the expressions (math\_errhandling & MATH\_ERRNO) and (math\_errhandling &  
4666 MATH\_ERREXCEPT) are independent of each other, but at least one of them must be non-zero.4667 **RATIONALE**

4668 None.

4669 **FUTURE DIRECTIONS**

4670 None.

4671 **SEE ALSO**

4672 *atan2*(), *feclearexcept*(), *fetestexcept*(), *isnan*(), *tan*(), the Base Definitions volume of |  
 4673 IEEE Std 1003.1-200x, Section 4.18, Treatment of Error Conditions for Mathematical Functions, |  
 4674 <math.h>

4675 **CHANGE HISTORY**

4676 First released in Issue 1. Derived from Issue 1 of the SVID.

4677 **Issue 5**

4678 The DESCRIPTION is updated to indicate how an application should check for an error. This  
4679 text was previously published in the APPLICATION USAGE section.

4680 **Issue 6**

4681 The *atanf()* and *atanl()* functions are added for alignment with the ISO/IEC 9899:1999 standard.

4682 The DESCRIPTION, RETURN VALUE, ERRORS, and APPLICATION USAGE sections are  
4683 revised to align with the ISO/IEC 9899:1999 standard.

4684 IEC 60559:1989 standard floating-point extensions over the ISO/IEC 9899:1999 standard are  
4685 marked.

4686 **NAME**

4687         atan2, atan2f, atan2l — arc tangent functions

4688 **SYNOPSIS**

4689         #include &lt;math.h&gt;

4690         double atan2(double y, double x);

4691         float atan2f(float y, float x);

4692         long double atan2l(long double y, long double x);

4693 **DESCRIPTION**

4694 **CX**         The functionality described on this reference page is aligned with the ISO C standard. Any  
 4695 conflict between the requirements described here and the ISO C standard is unintentional. This  
 4696 volume of IEEE Std 1003.1-200x defers to the ISO C standard.

4697         These functions shall compute the principal value of the arc tangent of  $y/x$ , using the signs of  
 4698 both arguments to determine the quadrant of the return value.

4699         An application wishing to check for error situations should set *errno* to zero and call  
 4700 *feclearexcept*(FE\_ALL\_EXCEPT) before calling these functions. On return, if *errno* is non-zero or  
 4701 *fetestexcept*(FE\_INVALID | FE\_DIVBYZERO | FE\_OVERFLOW | FE\_UNDERFLOW) is non-  
 4702 zero, an error has occurred.

4703 **RETURN VALUE**

4704         Upon successful completion, these functions shall return the arc tangent of  $y/x$  in the range  
 4705  $[-\pi, \pi]$  radians.

4706         If  $y$  is  $\pm 0$  and  $x$  is  $< 0$ ,  $\pm\pi$  shall be returned.

4707         If  $y$  is  $\pm 0$  and  $x$  is  $> 0$ ,  $\pm 0$  shall be returned.

4708         If  $y$  is  $< 0$  and  $x$  is  $\pm 0$ ,  $-\pi/2$  shall be returned.

4709         If  $y$  is  $> 0$  and  $x$  is  $\pm 0$ ,  $\pi/2$  shall be returned.

4710         If  $x$  is 0, a pole error shall not occur.

4711 **MX**         If either  $x$  or  $y$  is NaN, a NaN shall be returned.

4712         If the result underflows, a range error may occur and  $y/x$  should be returned.

4713         If  $y$  is  $\pm 0$  and  $x$  is  $-0$ ,  $\pm\pi$  shall be returned.

4714         If  $y$  is  $\pm 0$  and  $x$  is  $+0$ ,  $\pm 0$  shall be returned.

4715         For finite values of  $\pm y > 0$ , if  $x$  is  $-\text{Inf}$ ,  $\pm\pi$  shall be returned.

4716         For finite values of  $\pm y > 0$ , if  $x$  is  $+\text{Inf}$ ,  $\pm 0$  shall be returned.

4717         For finite values of  $x$ , if  $y$  is  $\pm\text{Inf}$ ,  $\pm\pi/2$  shall be returned.

4718         If  $y$  is  $\pm\text{Inf}$  and  $x$  is  $-\text{Inf}$ ,  $\pm 3\pi/4$  shall be returned.

4719         If  $y$  is  $\pm\text{Inf}$  and  $x$  is  $+\text{Inf}$ ,  $\pm\pi/4$  shall be returned.

4720         If both arguments are 0, a domain error shall not occur.

4721 **ERRORS**

4722         These functions may fail if:

4723 **MX**         **Range Error**         The result underflows.

4724         If the integer expression (math\_errhandling & MATH\_ERRNO) is non-zero, |  
 4725 then *errno* shall be set to [ERANGE]. If the integer expression |

4726 (math\_errhandling & MATH\_ERREXCEPT) is non-zero, then the underflow |  
4727 floating-point exception shall be raised. |

4728 **EXAMPLES**

4729 None.

4730 **APPLICATION USAGE**

4731 On error, the expressions (math\_errhandling & MATH\_ERRNO) and (math\_errhandling &  
4732 MATH\_ERREXCEPT) are independent of each other, but at least one of them must be non-zero.

4733 **RATIONALE**

4734 None.

4735 **FUTURE DIRECTIONS**

4736 None.

4737 **SEE ALSO**

4738 *atan()*, *feclearexcept()*, *fetestexcept()*, *isnan()*, *tan()*, the Base Definitions volume of |  
4739 IEEE Std 1003.1-200x, Section 4.18, Treatment of Error Conditions for Mathematical Functions, |  
4740 <math.h>

4741 **CHANGE HISTORY**

4742 First released in Issue 1. Derived from Issue 1 of the SVID.

4743 **Issue 5**

4744 The DESCRIPTION is updated to indicate how an application should check for an error. This  
4745 text was previously published in the APPLICATION USAGE section.

4746 **Issue 6**

4747 The *atan2f()* and *atan2l()* functions are added for alignment with the ISO/IEC 9899:1999  
4748 standard.

4749 The DESCRIPTION, RETURN VALUE, ERRORS, and APPLICATION USAGE sections are  
4750 revised to align with the ISO/IEC 9899:1999 standard.

4751 IEC 60559:1989 standard floating-point extensions over the ISO/IEC 9899:1999 standard are  
4752 marked.

4753 **NAME**4754        **atanf** — arc tangent function4755 **SYNOPSIS**

4756        #include &lt;math.h&gt;

4757        float atanf(float x);

4758 **DESCRIPTION**4759        Refer to *atan()*.



4760 **NAME**

4761           atanh, atanhf, atanh1 — inverse hyperbolic tangent functions

4762 **SYNOPSIS**

4763           #include &lt;math.h&gt;

4764           double atanh(double x);

4765           float atanhf(float x);

4766           long double atanh1(long double x);

4767 **DESCRIPTION**4768 **CX**       The functionality described on this reference page is aligned with the ISO C standard. Any  
4769 conflict between the requirements described here and the ISO C standard is unintentional. This  
4770 volume of IEEE Std 1003.1-200x defers to the ISO C standard.4771       These functions shall compute the inverse hyperbolic tangent of their argument *x*.4772       An application wishing to check for error situations should set *errno* to zero and call  
4773 *feclearexcept*(FE\_ALL\_EXCEPT) before calling these functions. On return, if *errno* is non-zero or  
4774 *fetestexcept*(FE\_INVALID | FE\_DIVBYZERO | FE\_OVERFLOW | FE\_UNDERFLOW) is non-  
4775 zero, an error has occurred.4776 **RETURN VALUE**4777       Upon successful completion, these functions shall return the inverse hyperbolic tangent of their  
4778 argument.4779       If *x* is  $\pm 1$ , a pole error shall occur, and *atanh()*, *atanhf()*, and *atanh1()* shall return the value of the  
4780 macro HUGE\_VAL, HUGE\_VALF, and HUGE\_VALL, respectively, with the same sign as the  
4781 correct value of the function.4782 **MX**       For finite  $|x| > 1$ , a domain error shall occur, and either a NaN (if supported), or an  
4783 implementation-defined value shall be returned.4784 **MX**       If *x* is NaN, a NaN shall be returned.4785       If *x* is  $\pm 0$ , *x* shall be returned.4786       If *x* is  $\pm \text{Inf}$ , a domain error shall occur, and either a NaN (if supported), or an implementation-  
4787 defined value shall be returned.4788       If *x* is subnormal, a range error may occur and *x* should be returned.4789 **ERRORS**

4790       These functions shall fail if:

4791 **MX**       Domain Error    The *x* argument is finite and not in the range  $[-1, 1]$ , or is  $\pm \text{Inf}$ .4792           If the integer expression (math\_errhandling & MATH\_ERRNO) is non-zero, |  
4793 then *errno* shall be set to [EDOM]. If the integer expression (math\_errhandling |  
4794 & MATH\_ERREXCEPT) is non-zero, then the invalid floating-point exception |  
4795 shall be raised. |4796       Pole Error        The *x* argument is  $\pm 1$ .4797           If the integer expression (math\_errhandling & MATH\_ERRNO) is non-zero, |  
4798 then *errno* shall be set to [ERANGE]. If the integer expression |  
4799 (math\_errhandling & MATH\_ERREXCEPT) is non-zero, then the divide-by- |  
4800 zero floating-point exception shall be raised. |

4801       These functions may fail if:

4802 MX     **Range Error**     The value of *x* is subnormal.

4803     If the integer expression (`math_errhandling & MATH_ERRNO`) is non-zero, |  
4804     then *errno* shall be set to [ERANGE]. If the integer expression |  
4805     (`math_errhandling & MATH_ERREXCEPT`) is non-zero, then the underflow |  
4806     floating-point exception shall be raised. |

4807 **EXAMPLES**

4808     None.

4809 **APPLICATION USAGE**

4810     On error, the expressions (`math_errhandling & MATH_ERRNO`) and (`math_errhandling &`  
4811     `MATH_ERREXCEPT`) are independent of each other, but at least one of them must be non-zero.

4812 **RATIONALE**

4813     None.

4814 **FUTURE DIRECTIONS**

4815     None.

4816 **SEE ALSO**

4817     *feclearexcept()*, *fetestexcept()*, *tanh()*, the Base Definitions volume of IEEE Std 1003.1-200x, Section |  
4818     4.18, Treatment of Error Conditions for Mathematical Functions, <**math.h**> |

4819 **CHANGE HISTORY**

4820     First released in Issue 4, Version 2.

4821 **Issue 5**

4822     Moved from X/OPEN UNIX extension to BASE.

4823 **Issue 6**

4824     The *atanh()* function is no longer marked as an extension.

4825     The *atanhf()*, and *atanhl()* functions are added for alignment with the ISO/IEC 9899:1999  
4826     standard.

4827     The DESCRIPTION, RETURN VALUE, ERRORS, and APPLICATION USAGE sections are  
4828     revised to align with the ISO/IEC 9899:1999 standard.

4829     IEC 60559:1989 standard floating-point extensions over the ISO/IEC 9899:1999 standard are  
4830     marked.

4831 **NAME**

4832       atanl — arc tangent function

4833 **SYNOPSIS**

4834       #include <math.h>

4835       long double atanl(long double x);

4836 **DESCRIPTION**

4837       Refer to *atan()*.

4838 **NAME**

4839         atexit — register a function to run at process termination

4840 **SYNOPSIS**

4841         #include &lt;stdlib.h&gt;

4842         int atexit(void (\*func)(void));

4843 **DESCRIPTION**4844 **CX**         The functionality described on this reference page is aligned with the ISO C standard. Any  
4845 conflict between the requirements described here and the ISO C standard is unintentional. This  
4846 volume of IEEE Std 1003.1-200x defers to the ISO C standard.4847         The *atexit()* function shall register the function pointed to by *func*, to be called without  
4848 arguments at normal program termination. At normal program termination, all functions  
4849 registered by the *atexit()* function shall be called, in the reverse order of their registration, except  
4850 that a function is called after any previously registered functions that had already been called at  
4851 the time it was registered. Normal termination occurs either by a call to *exit()* or a return from  
4852 *main()*.4853         At least 32 functions can be registered with *atexit()*.4854 **CX**         After a successful call to any of the *exec* functions, any functions previously registered by *atexit()*  
4855 shall no longer be registered.4856 **RETURN VALUE**4857         Upon successful completion, *atexit()* shall return 0; otherwise, it shall return a non-zero value.4858 **ERRORS**

4859         No errors are defined.

4860 **EXAMPLES**

4861         None.

4862 **APPLICATION USAGE**4863         The functions registered by a call to *atexit()* must return to ensure that all registered functions  
4864 are called.4865         The application should call *sysconf()* to obtain the value of {ATEXIT\_MAX}, the number of  
4866 functions that can be registered. There is no way for an application to tell how many functions  
4867 have already been registered with *atexit()*.4868 **RATIONALE**

4869         None.

4870 **FUTURE DIRECTIONS**

4871         None.

4872 **SEE ALSO**4873         *exit()*, *sysconf()*, the Base Definitions volume of IEEE Std 1003.1-200x, <stdlib.h>4874 **CHANGE HISTORY**

4875         First released in Issue 4. Derived from the ANSI C standard.

4876 **Issue 6**

4877         Extensions beyond the ISO C standard are now marked.

4878         The DESCRIPTION is updated for alignment with the ISO/IEC 9899:1999 standard.

4879 **NAME**

4880            atof — convert a string to double-precision number

4881 **SYNOPSIS**

4882            #include &lt;stdlib.h&gt;

4883            double atof(const char \*str);

4884 **DESCRIPTION**

4885 cx        The functionality described on this reference page is aligned with the ISO C standard. Any  
4886        conflict between the requirements described here and the ISO C standard is unintentional. This  
4887        volume of IEEE Std 1003.1-200x defers to the ISO C standard.

4888        The call *atof(str)* shall be equivalent to:

4889            strtod(str, (char \*\*)NULL),

4890        except that the handling of errors may differ. If the value cannot be represented, the behavior is  
4891        undefined.

4892 **RETURN VALUE**4893        The *atof()* function shall return the converted value if the value can be represented.4894 **ERRORS**

4895        No errors are defined.

4896 **EXAMPLES**

4897        None.

4898 **APPLICATION USAGE**

4899        The *atof()* function is subsumed by *strtod()* but is retained because it is used extensively in  
4900        existing code. If the number is not known to be in range, *strtod()* should be used because *atof()* is  
4901        not required to perform any error checking.

4902 **RATIONALE**

4903        None.

4904 **FUTURE DIRECTIONS**

4905        None.

4906 **SEE ALSO**4907        *strtod()*, the Base Definitions volume of IEEE Std 1003.1-200x, <stdlib.h>4908 **CHANGE HISTORY**

4909        First released in Issue 1. Derived from Issue 1 of the SVID.

4910 **NAME**

4911         atoi — convert a string to an integer

4912 **SYNOPSIS**

4913         #include &lt;stdlib.h&gt;

4914         int atoi(const char \*str);

4915 **DESCRIPTION**

4916 cx         The functionality described on this reference page is aligned with the ISO C standard. Any  
4917         conflict between the requirements described here and the ISO C standard is unintentional. This  
4918         volume of IEEE Std 1003.1-200x defers to the ISO C standard.

4919         The call *atoi(str)* shall be equivalent to:

4920         (int) strtol(str, (char \*\*)NULL, 10)

4921         except that the handling of errors may differ. If the value cannot be represented, the behavior is  
4922         undefined.

4923 **RETURN VALUE**4924         The *atoi()* function shall return the converted value if the value can be represented.4925 **ERRORS**

4926         No errors are defined.

4927 **EXAMPLES**4928         **Converting an Argument**

4929         The following example checks for proper usage of the program. If there is an argument and the  
4930         decimal conversion of this argument (obtained using *atoi()*) is greater than 0, then the program  
4931         has a valid number of minutes to wait for an event.

```
4932         #include <stdlib.h>
4933         #include <stdio.h>
4934         ...
4935         int minutes_to_event;
4936         ...
4937         if (argc < 2 || ((minutes_to_event = atoi (argv[1]))) <= 0) {
4938             fprintf(stderr, "Usage: %s minutes\n", argv[0]); exit(1);
4939         }
4940         ...
```

4941 **APPLICATION USAGE**

4942         The *atoi()* function is subsumed by *strtol()* but is retained because it is used extensively in  
4943         existing code. If the number is not known to be in range, *strtol()* should be used because *atoi()* is  
4944         not required to perform any error checking.

4945 **RATIONALE**

4946         None.

4947 **FUTURE DIRECTIONS**

4948         None.

4949 **SEE ALSO**4950         *strtol()*, the Base Definitions volume of IEEE Std 1003.1-200x, <stdlib.h>

4951 **CHANGE HISTORY**

4952 First released in Issue 1. Derived from Issue 1 of the SVID.

4953 **NAME**4954 `atol`, `atoll` — convert a string to a long integer4955 **SYNOPSIS**4956 `#include <stdlib.h>`4957 `long atol(const char *str);`4958 `long long atoll(const char *nptr);`4959 **DESCRIPTION**

4960 **cx** The functionality described on this reference page is aligned with the ISO C standard. Any  
4961 conflict between the requirements described here and the ISO C standard is unintentional. This  
4962 volume of IEEE Std 1003.1-200x defers to the ISO C standard.

4963 The call `atol(str)` shall be equivalent to:4964 `strtoul(str, (char **)NULL, 10)`4965 The call `atoll(str)` shall be equivalent to:4966 `strtoll(nptr, (char **)NULL, 10)`

4967 except that the handling of errors may differ. If the value cannot be represented, the behavior is  
4968 undefined.

4969 **RETURN VALUE**

4970 These functions shall return the converted value if the value can be represented.

4971 **ERRORS**

4972 No errors are defined.

4973 **EXAMPLES**

4974 None.

4975 **APPLICATION USAGE**

4976 The `atol()` function is subsumed by `strtoul()` but is retained because it is used extensively in  
4977 existing code. If the number is not known to be in range, `strtoul()` should be used because `atol()` is  
4978 not required to perform any error checking.

4979 **RATIONALE**

4980 None.

4981 **FUTURE DIRECTIONS**

4982 None.

4983 **SEE ALSO**4984 `strtoul()`, the Base Definitions volume of IEEE Std 1003.1-200x, `<stdlib.h>`4985 **CHANGE HISTORY**

4986 First released in Issue 1. Derived from Issue 1 of the SVID.

4987 **Issue 6**4988 The `atoll()` function is added for alignment with the ISO/IEC 9899:1999 standard.



4989 **NAME**4990 `basename` — return the last component of a pathname |4991 **SYNOPSIS**4992 `XSI` `#include <libgen.h>`4993 `char *basename(char *path);`

4994

4995 **DESCRIPTION**4996 The `basename()` function shall take the pathname pointed to by `path` and return a pointer to the |  
4997 final component of the pathname, deleting any trailing `'/'` characters. |4998 If the string consists entirely of the `'/'` character, `basename()` shall return a pointer to the string  
4999 `"/"`. If the string is exactly `"/"`, it is implementation-defined whether `'/'` or `"/"` is  
5000 returned.5001 If `path` is a null pointer or points to an empty string, `basename()` shall return a pointer to the  
5002 string  `"."`.5003 The `basename()` function may modify the string pointed to by `path`, and may return a pointer to  
5004 static storage that may then be overwritten by a subsequent call to `basename()`.5005 The `basename()` function need not be reentrant. A function that is not required to be reentrant is  
5006 not required to be thread-safe.5007 **RETURN VALUE**5008 The `basename()` function shall return a pointer to the final component of `path`.5009 **ERRORS**

5010 No errors are defined.

5011 **EXAMPLES**5012 **Using `basename()`**5013 The following program fragment returns a pointer to the value `lib`, which is the base name of  
5014 `/usr/lib`.5015 `#include <libgen.h>`  
5016 `...`  
5017 `char *name = "/usr/lib";`  
5018 `char *base;`  
5019 `base = basename(name);`  
5020 `...`5021 **Sample Input and Output Strings for `basename()`**5022 In the following table, the input string is the value pointed to by `path`, and the output string is  
5023 the return value of the `basename()` function.

5024

5025

5026

5027

Input String	Output String
<code>"/usr/lib"</code>	<code>"lib"</code>
<code>"/usr/"</code>	<code>"usr"</code>
<code>"/"</code>	<code>"/"</code>

5028 **APPLICATION USAGE**

5029       None.

5030 **RATIONALE**

5031       None.

5032 **FUTURE DIRECTIONS**

5033       None.

5034 **SEE ALSO**5035       *dirname()*, the Base Definitions volume of IEEE Std 1003.1-200x, <**libgen.h**>, the Shell and  
5036       Utilities volume of IEEE Std 1003.1-200x, *basename*5037 **CHANGE HISTORY**

5038       First released in Issue 4, Version 2.

5039 **Issue 5**

5040       Moved from X/OPEN UNIX extension to BASE.

5041       Normative text previously in the APPLICATION USAGE section is moved to the  
5042       DESCRIPTION.

5043       A note indicating that this function need not be reentrant is added to the DESCRIPTION.

5044 **Issue 6**

5045       In the DESCRIPTION, the note about reentrancy is expanded to cover thread-safety.

5046 **NAME**5047       bcmp — memory operations (**LEGACY**)5048 **SYNOPSIS**

5049 XSI       #include &lt;strings.h&gt;

5050       int bcmp(const void \*s1, const void \*s2, size\_t n);

5051

5052 **DESCRIPTION**5053       The *bcmp()* function shall compare the first *n* bytes of the area pointed to by *s1* with the area  
5054       pointed to by *s2*.5055 **RETURN VALUE**5056       The *bcmp()* function shall return 0 if *s1* and *s2* are identical; otherwise, it shall return non-zero.  
5057       Both areas are assumed to be *n* bytes long. If the value of *n* is 0, *bcmp()* shall return 0.5058 **ERRORS**

5059       No errors are defined.

5060 **EXAMPLES**

5061       None.

5062 **APPLICATION USAGE**5063       *memcmp()* is preferred over this function.5064       For maximum portability, it is recommended to replace the function call to *bcmp()* as follows:

5065       #define bcmp(b1,b2,len) memcmp((b1), (b2), (size\_t)(len))

5066 **RATIONALE**

5067       None.

5068 **FUTURE DIRECTIONS**

5069       This function may be withdrawn in a future version.

5070 **SEE ALSO**5071       *memcmp()*, the Base Definitions volume of IEEE Std 1003.1-200x, <strings.h>5072 **CHANGE HISTORY**

5073       First released in Issue 4, Version 2.

5074 **Issue 5**

5075       Moved from X/OPEN UNIX extension to BASE.

5076 **Issue 6**

5077       This function is marked LEGACY.

5078 **NAME**5079 **bcopy** — memory operations (**LEGACY**)5080 **SYNOPSIS**5081 XSI 

```
#include <strings.h>
```

5082 

```
void bcopy(const void *s1, void *s2, size_t n);
```

5083

5084 **DESCRIPTION**5085 The *bcopy()* function shall copy *n* bytes from the area pointed to by *s1* to the area pointed to by  
5086 *s2*.5087 The bytes are copied correctly even if the area pointed to by *s1* overlaps the area pointed to by  
5088 *s2*.5089 **RETURN VALUE**5090 The *bcopy()* function shall not return a value.5091 **ERRORS**

5092 No errors are defined.

5093 **EXAMPLES**

5094 None.

5095 **APPLICATION USAGE**5096 *memmove()* is preferred over this function.

5097 The following are approximately equivalent (note the order of the arguments):

5098 

```
bcopy(s1,s2,n) ~ memmove(s2,s1,n)
```

5099 For maximum portability, it is recommended to replace the function call to *bcopy()* as follows:5100 

```
#define bcopy(b1,b2,len) (memmove((b2), (b1), (len)), (void) 0)
```

5101 **RATIONALE**

5102 None.

5103 **FUTURE DIRECTIONS**

5104 This function may be withdrawn in a future version.

5105 **SEE ALSO**5106 *memmove()*, the Base Definitions volume of IEEE Std 1003.1-200x, <**strings.h**>5107 **CHANGE HISTORY**

5108 First released in Issue 4, Version 2.

5109 **Issue 5**

5110 Moved from X/OPEN UNIX extension to BASE.

5111 **Issue 6**

5112 This function is marked LEGACY.

5113 **NAME**

5114        bind — bind a name to a socket

5115 **SYNOPSIS**

5116        #include &lt;sys/socket.h&gt;

5117        int bind(int *socket*, const struct sockaddr \**address*,  
5118                socklen\_t *address\_len*);5119 **DESCRIPTION**5120        The *bind()* function shall assign a local socket address *address* to a socket identified by descriptor  
5121        *socket* that has no local socket address assigned. Sockets created with the *socket()* function are  
5122        initially unnamed; they are identified only by their address family.5123        The *bind()* function takes the following arguments:5124        *socket*                Specifies the file descriptor of the socket to be bound.5125        *address*               Points to a **sockaddr** structure containing the address to be bound to the  
5126        socket. The length and format of the address depend on the address family of  
5127        the socket.5128        *address\_len*         Specifies the length of the **sockaddr** structure pointed to by the *address*  
5129        argument.5130        The socket specified by *socket* may require the process to have appropriate privileges to use the  
5131        *bind()* function.5132 **RETURN VALUE**5133        Upon successful completion, *bind()* shall return 0; otherwise,  $-1$  shall be returned and *errno* set  
5134        to indicate the error.5135 **ERRORS**5136        The *bind()* function shall fail if:

5137        [EADDRINUSE]

5138                The specified address is already in use.

5139        [EADDRNOTAVAIL]

5140                The specified address is not available from the local machine.

5141        [EAFNOSUPPORT]

5142                The specified address is not a valid address for the address family of the  
5143        specified socket.5144        [EBADF]                The *socket* argument is not a valid file descriptor.5145        [EINVAL]               The socket is already bound to an address, and the protocol does not support  
5146        binding to a new address; or the socket has been shut down.5147        [ENOTSOCK]            The *socket* argument does not refer to a socket.5148        [EOPNOTSUPP]         The socket type of the specified socket does not support binding to an  
5149        address.5150        If the address family of the socket is AF\_UNIX, then *bind()* shall fail if:5151        [EACCES]               A component of the path prefix denies search permission, or the requested  
5152        name requires writing in a directory with a mode that denies write  
5153        permission.

5154	[EDESTADDRREQ] or [EISDIR]	
5155		The <i>address</i> argument is a null pointer.
5156	[EIO]	An I/O error occurred.
5157	[ELOOP]	A loop exists in symbolic links encountered during resolution of the pathname
5158		in <i>address</i> .
5159	[ENAMETOOLONG]	
5160		A component of a pathname exceeded {NAME_MAX} characters, or an entire
5161		pathname exceeded {PATH_MAX} characters.
5162	[ENOENT]	A component of the pathname does not name an existing file or the pathname
5163		is an empty string.
5164	[ENOTDIR]	A component of the path prefix of the pathname in <i>address</i> is not a directory.
5165	[EROFS]	The name would reside on a read-only file system.
5166		The <i>bind()</i> function may fail if:
5167	[EACCES]	The specified address is protected and the current user does not have
5168		permission to bind to it.
5169	[EINVAL]	The <i>address_len</i> argument is not a valid length for the address family.
5170	[EISCONN]	The socket is already connected.
5171	[ELOOP]	More than {SYMLOOP_MAX} symbolic links were encountered during
5172		resolution of the pathname in <i>address</i> .
5173	[ENAMETOOLONG]	
5174		Pathname resolution of a symbolic link produced an intermediate result
5175		whose length exceeds {PATH_MAX}.
5176	[ENOBUFS]	Insufficient resources were available to complete the call.
5177	<b>EXAMPLES</b>	
5178		None.
5179	<b>APPLICATION USAGE</b>	
5180		An application program can retrieve the assigned socket name with the <i>getsockname()</i> function.
5181	<b>RATIONALE</b>	
5182		None.
5183	<b>FUTURE DIRECTIONS</b>	
5184		None.
5185	<b>SEE ALSO</b>	
5186		<i>connect()</i> , <i>getsockname()</i> , <i>listen()</i> , <i>socket()</i> , the Base Definitions volume of IEEE Std 1003.1-200x,
5187		<sys/socket.h>
5188	<b>CHANGE HISTORY</b>	
5189		First released in Issue 6. Derived from the XNS, Issue 5.2 specification.

5190 **NAME**

5191        bsd\_signal — simplified signal facilities

5192 **SYNOPSIS**

5193 OB XSI   #include &lt;signal.h&gt;

5194        void (\*bsd\_signal(int sig, void (\*func)(int)))(int);

5195

5196 **DESCRIPTION**5197        The *bsd\_signal()* function provides a partially compatible interface for programs written to  
5198        historical system interfaces (see APPLICATION USAGE).5199        The function call *bsd\_signal(sig, func)* shall be equivalent to the following:

5200        void (\*bsd\_signal(int sig, void (\*func)(int)))(int)

```

5201        {
5202            struct sigaction act, oact;
5203            act.sa_handler = func;
5204            act.sa_flags = SA_RESTART;
5205            sigemptyset(&act.sa_mask);
5206            sigaddset(&act.sa_mask, sig);
5207            if (sigaction(sig, &act, &oact) == -1)
5208                return(SIG_ERR);
5209            return(oact.sa_handler);
5210        }

```

5211        The handler function should be declared:

5212        void handler(int sig);

5213        where *sig* is the signal number. The behavior is undefined if *func* is a function that takes more  
5214        than one argument, or an argument of a different type.5215 **RETURN VALUE**5216        Upon successful completion, *bsd\_signal()* shall return the previous action for *sig*. Otherwise,  
5217        SIG\_ERR shall be returned and *errno* shall be set to indicate the error.5218 **ERRORS**5219        Refer to *sigaction()*.5220 **EXAMPLES**

5221        None.

5222 **APPLICATION USAGE**5223        This function is a direct replacement for the BSD *signal()* function for simple applications that  
5224        are installing a single-argument signal handler function. If a BSD signal handler function is being  
5225        installed that expects more than one argument, the application has to be modified to use  
5226        *sigaction()*. The *bsd\_signal()* function differs from *signal()* in that the SA\_RESTART flag is set  
5227        and the SA\_RESETHAND is clear when *bsd\_signal()* is used. The state of these flags is not  
5228        specified for *signal()*.5229        It is recommended that new applications use the *sigaction()* function.5230 **RATIONALE**

5231        None.

5232 **FUTURE DIRECTIONS**

5233 None.

5234 **SEE ALSO**5235 *sigaction()*, *sigaddset()*, *sigemptyset()*, *signal()*, the Base Definitions volume of  
5236 IEEE Std 1003.1-200x, <**signal.h**>5237 **CHANGE HISTORY**

5238 First released in Issue 4, Version 2.

5239 **Issue 5**

5240 Moved from X/OPEN UNIX extension to BASE.

5241 **Issue 6**

5242 This function is marked obsolescent.



5243 **NAME**5244       **bsearch** — binary search a sorted table5245 **SYNOPSIS**

5246       #include &lt;stdlib.h&gt;

5247       void \*bsearch(const void \*key, const void \*base, size\_t nel,  
5248                   size\_t width, int (\*compar)(const void \*, const void \*));5249 **DESCRIPTION**5250 **CX**     The functionality described on this reference page is aligned with the ISO C standard. Any  
5251     conflict between the requirements described here and the ISO C standard is unintentional. This  
5252     volume of IEEE Std 1003.1-200x defers to the ISO C standard.5253     The *bsearch()* function shall search an array of *nel* objects, the initial element of which is pointed  
5254     to by *base*, for an element that matches the object pointed to by *key*. The size of each element in  
5255     the array is specified by *width*.5256     The comparison function pointed to by *compar* shall be called with two arguments that point to  
5257     the *key* object and to an array element, in that order.5258     The application shall ensure that the function returns an integer less than, equal to, or greater  
5259     than 0 if the *key* object is considered, respectively, to be less than, to match, or to be greater than  
5260     the array element. The application shall ensure that the array consists of all the elements that  
5261     compare less than, all the elements that compare equal to, and all the elements that compare  
5262     greater than the *key* object, in that order.5263 **RETURN VALUE**5264     The *bsearch()* function shall return a pointer to a matching member of the array, or a null pointer  
5265     if no match is found. If two or more members compare equal, which member is returned is  
5266     unspecified.5267 **ERRORS**

5268     No errors are defined.

5269 **EXAMPLES**5270     The example below searches a table containing pointers to nodes consisting of a string and its  
5271     length. The table is ordered alphabetically on the string in the node pointed to by each entry.5272     The code fragment below reads in strings and either finds the corresponding node and prints out  
5273     the string and its length, or prints an error message.5274     #include <stdio.h>  
5275     #include <stdlib.h>  
5276     #include <string.h>  
  
5277     #define TABSIZE     1000  
  
5278     struct node {   /\* These are stored in the table. \*/  
5279         char \*string;  
5280         int length;  
5281     };  
5282     struct node table[TABSIZE];     /\* Table to be searched. \*/  
5283     .  
5284     .  
5285     .  
5286     {  
5287         struct node \*node\_ptr, node;  
5288         /\* routine to compare 2 nodes \*/

```

5289     int node_compare(const void *, const void *);
5290     char str_space[20]; /* Space to read string into. */
5291     .
5292     .
5293     .
5294     node.string = str_space;
5295     while (scanf("%s", node.string) != EOF) {
5296         node_ptr = (struct node *)bsearch((void *)&node,
5297             (void *)table, TABSIZE,
5298             sizeof(struct node), node_compare);
5299         if (node_ptr != NULL) {
5300             (void)printf("string = %20s, length = %d\n",
5301                 node_ptr->string, node_ptr->length);
5302         } else {
5303             (void)printf("not found: %s\n", node.string);
5304         }
5305     }
5306 }
5307 /*
5308     This routine compares two nodes based on an
5309     alphabetical ordering of the string field.
5310 */
5311 int
5312 node_compare(const void *node1, const void *node2)
5313 {
5314     return strcoll(((const struct node *)node1)->string,
5315         ((const struct node *)node2)->string);
5316 }

```

**5317 APPLICATION USAGE**

5318 The pointers to the key and the element at the base of the table should be of type pointer-to-  
5319 element.

5320 The comparison function need not compare every byte, so arbitrary data may be contained in  
5321 the elements in addition to the values being compared.

5322 In practice, the array is usually sorted according to the comparison function.

**5323 RATIONALE**

5324 None.

**5325 FUTURE DIRECTIONS**

5326 None.

**5327 SEE ALSO**

5328 *hcreate()*, *lsearch()*, *qsort()*, *tsearch()*, the Base Definitions volume of IEEE Std 1003.1-200x,  
5329 <stdlib.h>

**5330 CHANGE HISTORY**

5331 First released in Issue 1. Derived from Issue 1 of the SVID.

**5332 Issue 6**

5333 The DESCRIPTION is updated to avoid use of the term “must” for application requirements.

5334 **NAME**

5335       btowc — single byte to wide character conversion

5336 **SYNOPSIS**

5337       #include &lt;stdio.h&gt;

5338       #include &lt;wchar.h&gt;

5339       wint\_t btowc(int c);

5340 **DESCRIPTION**

5341 cx     The functionality described on this reference page is aligned with the ISO C standard. Any  
5342     conflict between the requirements described here and the ISO C standard is unintentional. This  
5343     volume of IEEE Std 1003.1-200x defers to the ISO C standard.

5344     The *btowc()* function shall determine whether *c* constitutes a valid (one-byte) character in the  
5345     initial shift state.

5346     The behavior of this function shall be affected by the *LC\_CTYPE* category of the current locale.

5347 **RETURN VALUE**

5348     The *btowc()* function shall return WEOF if *c* has the value EOF or if (**unsigned char**) *c* does not  
5349     constitute a valid (one-byte) character in the initial shift state. Otherwise, it shall return the  
5350     wide-character representation of that character.

5351 **ERRORS**

5352       No errors are defined.

5353 **EXAMPLES**

5354       None.

5355 **APPLICATION USAGE**

5356       None.

5357 **RATIONALE**

5358       None.

5359 **FUTURE DIRECTIONS**

5360       None.

5361 **SEE ALSO**5362       *wctob()*, the Base Definitions volume of IEEE Std 1003.1-200x, <wchar.h>5363 **CHANGE HISTORY**

5364       First released in Issue 5. Included for alignment with ISO/IEC 9899:1990/Amendment 1:1995  
5365       (E).

5366 **NAME**5367       **bzero** — memory operations (**LEGACY**)5368 **SYNOPSIS**

5369 XSI       #include &lt;strings.h&gt;

5370       void bzero(void \*s, size\_t n);

5371

5372 **DESCRIPTION**5373       The *bzero()* function shall place *n* zero-valued bytes in the area pointed to by *s*.5374 **RETURN VALUE**5375       The *bzero()* function shall not return a value.5376 **ERRORS**

5377       No errors are defined.

5378 **EXAMPLES**

5379       None.

5380 **APPLICATION USAGE**5381       *memset()* is preferred over this function.5382       For maximum portability, it is recommended to replace the function call to *bzero()* as follows:

5383       #define bzero(b,len) (memset((b), '\0', (len)), (void) 0)

5384 **RATIONALE**

5385       None.

5386 **FUTURE DIRECTIONS**

5387       This function may be withdrawn in a future version.

5388 **SEE ALSO**5389       *memset()*, the Base Definitions volume of IEEE Std 1003.1-200x, <**strings.h**>5390 **CHANGE HISTORY**

5391       First released in Issue 4, Version 2.

5392 **Issue 5**

5393       Moved from X/OPEN UNIX extension to BASE.

5394 **Issue 6**

5395       This function is marked LEGACY.

5396 **NAME**

5397       cabs, cabsf, cabsl — return a complex absolute value

5398 **SYNOPSIS**

5399       #include &lt;complex.h&gt;

5400       double cabs(double complex *z*);5401       float cabsf(float complex *z*);5402       long double cabsl(long double complex *z*);5403 **DESCRIPTION**5404 **CX**       The functionality described on this reference page is aligned with the ISO C standard. Any  
5405 conflict between the requirements described here and the ISO C standard is unintentional. This  
5406 volume of IEEE Std 1003.1-200x defers to the ISO C standard.5407       These functions shall compute the complex absolute value (also called norm, modulus, or  
5408 magnitude) of *z*.5409 **RETURN VALUE**

5410       These functions shall return the complex absolute value.

5411 **ERRORS**

5412       No errors are defined.

5413 **EXAMPLES**

5414       None.

5415 **APPLICATION USAGE**

5416       None.

5417 **RATIONALE**

5418       None.

5419 **FUTURE DIRECTIONS**

5420       None.

5421 **SEE ALSO**

5422       The Base Definitions volume of IEEE Std 1003.1-200x, &lt;complex.h&gt;

5423 **CHANGE HISTORY**

5424       First released in Issue 6. Derived from the ISO/IEC 9899:1999 standard.

5425 **NAME**

5426 cacos, cacosf, cacosl — complex arc cosine functions

5427 **SYNOPSIS**

5428 #include &lt;complex.h&gt;

5429 double complex cacos(double complex z);

5430 float complex cacosf(float complex z);

5431 long double complex cacosl(long double complex z);

5432 **DESCRIPTION**

5433 cx The functionality described on this reference page is aligned with the ISO C standard. Any  
5434 conflict between the requirements described here and the ISO C standard is unintentional. This  
5435 volume of IEEE Std 1003.1-200x defers to the ISO C standard.

5436 These functions shall compute the complex arc cosine of  $z$ , with branch cuts outside the interval  
5437  $[-1, +1]$  along the real axis.

5438 **RETURN VALUE**

5439 These functions shall return the complex arc cosine value, in the range of a strip mathematically  
5440 unbounded along the imaginary axis and in the interval  $[0, \pi]$  along the real axis.

5441 **ERRORS**

5442 No errors are defined.

5443 **EXAMPLES**

5444 None.

5445 **APPLICATION USAGE**

5446 None.

5447 **RATIONALE**

5448 None.

5449 **FUTURE DIRECTIONS**

5450 None.

5451 **SEE ALSO**5452 *ccos()*, the Base Definitions volume of IEEE Std 1003.1-200x, <complex.h>5453 **CHANGE HISTORY**

5454 First released in Issue 6. Derived from the ISO/IEC 9899:1999 standard.

5455 **NAME**

5456        cacosf — complex arc cosine functions

5457 **SYNOPSIS**

5458        #include <complex.h>

5459        float complex cacosf(float complex z);

5460 **DESCRIPTION**

5461        Refer to *cacos()*.

5462 **NAME**

5463 cacosh, cacoshf, cacoshl — complex arc hyperbolic cosine functions

5464 **SYNOPSIS**

5465 #include &lt;complex.h&gt;

5466 double complex cacosh(double complex z);

5467 float complex cacoshf(float complex z);

5468 long double complex cacoshl(long double complex z);

5469 **DESCRIPTION**

5470 cx The functionality described on this reference page is aligned with the ISO C standard. Any  
5471 conflict between the requirements described here and the ISO C standard is unintentional. This  
5472 volume of IEEE Std 1003.1-200x defers to the ISO C standard.

5473 These functions shall compute the complex arc hyperbolic cosine of  $z$ , with a branch cut at  
5474 values less than 1 along the real axis.

5475 **RETURN VALUE**

5476 These functions shall return the complex arc hyperbolic cosine value, in the range of a half-strip  
5477 of non-negative values along the real axis and in the interval  $[-i\pi, +i\pi]$  along the imaginary axis.

5478 **ERRORS**

5479 No errors are defined.

5480 **EXAMPLES**

5481 None.

5482 **APPLICATION USAGE**

5483 None.

5484 **RATIONALE**

5485 None.

5486 **FUTURE DIRECTIONS**

5487 None.

5488 **SEE ALSO**5489 *ccosh()*, the Base Definitions volume of IEEE Std 1003.1-200x, <complex.h>5490 **CHANGE HISTORY**

5491 First released in Issue 6. Derived from the ISO/IEC 9899:1999 standard.



5492 **NAME**

5493       cacosl — complex arc cosine functions

5494 **SYNOPSIS**

5495       #include &lt;complex.h&gt;

5496       long double complex cacosl(long double complex z);

5497 **DESCRIPTION**5498       Refer to *acos()*.

5499 **NAME**

5500        calloc — a memory allocator

5501 **SYNOPSIS**

5502        #include &lt;stdlib.h&gt;

5503        void \*calloc(size\_t *nelem*, size\_t *elsize*);5504 **DESCRIPTION**

5505 **CX**        The functionality described on this reference page is aligned with the ISO C standard. Any  
5506 conflict between the requirements described here and the ISO C standard is unintentional. This  
5507 volume of IEEE Std 1003.1-200x defers to the ISO C standard.

5508        The *calloc()* function shall allocate unused space for an array of *nelem* elements each of whose  
5509 size in bytes is *elsize*. The space shall be initialized to all bits 0.

5510        The order and contiguity of storage allocated by successive calls to *calloc()* is unspecified. The  
5511 pointer returned if the allocation succeeds shall be suitably aligned so that it may be assigned to  
5512 a pointer to any type of object and then used to access such an object or an array of such objects  
5513 in the space allocated (until the space is explicitly freed or reallocated). Each such allocation  
5514 shall yield a pointer to an object disjoint from any other object. The pointer returned shall point  
5515 to the start (lowest byte address) of the allocated space. If the space cannot be allocated, a null  
5516 pointer shall be returned. If the size of the space requested is 0, the behavior is implementation-  
5517 defined: the value returned shall be either a null pointer or a unique pointer.

5518 **RETURN VALUE**

5519        Upon successful completion with both *nelem* and *elsize* non-zero, *calloc()* shall return a pointer to  
5520 the allocated space. If either *nelem* or *elsize* is 0, then either a null pointer or a unique pointer  
5521 value that can be successfully passed to *free()* shall be returned. Otherwise, it shall return a null  
5522 **CX** pointer and set *errno* to indicate the error.

5523 **ERRORS**5524        The *calloc()* function shall fail if:5525 **CX**        [ENOMEM]        Insufficient memory is available.5526 **EXAMPLES**

5527        None.

5528 **APPLICATION USAGE**

5529        There is now no requirement for the implementation to support the inclusion of &lt;malloc.h&gt;.

5530 **RATIONALE**

5531        None.

5532 **FUTURE DIRECTIONS**

5533        None.

5534 **SEE ALSO**5535        *free()*, *malloc()*, *realloc()*, the Base Definitions volume of IEEE Std 1003.1-200x, <stdlib.h>5536 **CHANGE HISTORY**

5537        First released in Issue 1. Derived from Issue 1 of the SVID.

5538 **Issue 6**

5539        Extensions beyond the ISO C standard are now marked.

5540        The following new requirements on POSIX implementations derive from alignment with the  
5541 Single UNIX Specification:

5542  
5543

- The setting of *errno* and the [ENOMEM] error condition are mandatory if an insufficient memory condition occurs.

5544 **NAME**5545 `carg`, `cargf`, `cargl` — complex argument functions5546 **SYNOPSIS**5547 `#include <complex.h>`5548 `double carg(double complex z);`5549 `float cargf(float complex z);`5550 `long double cargl(long double complex z);`5551 **DESCRIPTION**

5552 `CX` The functionality described on this reference page is aligned with the ISO C standard. Any  
5553 conflict between the requirements described here and the ISO C standard is unintentional. This  
5554 volume of IEEE Std 1003.1-200x defers to the ISO C standard.

5555 These functions shall compute the argument (also called phase angle) of  $z$ , with a branch cut  
5556 along the negative real axis.

5557 **RETURN VALUE**5558 These functions shall return the value of the argument in the interval  $[-\pi, +\pi]$ .5559 **ERRORS**

5560 No errors are defined.

5561 **EXAMPLES**

5562 None.

5563 **APPLICATION USAGE**

5564 None.

5565 **RATIONALE**

5566 None.

5567 **FUTURE DIRECTIONS**

5568 None.

5569 **SEE ALSO**5570 `cimag()`, `conj()`, `cproj()`, the Base Definitions volume of IEEE Std 1003.1-200x, `<complex.h>`5571 **CHANGE HISTORY**

5572 First released in Issue 6. Derived from the ISO/IEC 9899:1999 standard.

5573 **NAME**

5574            casin, casinf, casinl — complex arc sine functions

5575 **SYNOPSIS**

5576            #include &lt;complex.h&gt;

5577            double complex casin(double complex z);

5578            float complex casinf(float complex z);

5579            long double complex casinl(long double complex z);

5580 **DESCRIPTION**

5581 cx        The functionality described on this reference page is aligned with the ISO C standard. Any  
5582            conflict between the requirements described here and the ISO C standard is unintentional. This  
5583            volume of IEEE Std 1003.1-200x defers to the ISO C standard.

5584            These functions shall compute the complex arc sine of  $z$ , with branch cuts outside the interval  
5585             $[-1, +1]$  along the real axis.

5586 **RETURN VALUE**

5587            These functions shall return the complex arc sine value, in the range of a strip mathematically  
5588            unbounded along the imaginary axis and in the interval  $[-\pi/2, +\pi/2]$  along the real axis.

5589 **ERRORS**

5590            No errors are defined.

5591 **EXAMPLES**

5592            None.

5593 **APPLICATION USAGE**

5594            None.

5595 **RATIONALE**

5596            None.

5597 **FUTURE DIRECTIONS**

5598            None.

5599 **SEE ALSO**5600            `csin()`, the Base Definitions volume of IEEE Std 1003.1-200x, <complex.h>5601 **CHANGE HISTORY**

5602            First released in Issue 6. Derived from the ISO/IEC 9899:1999 standard.

5603 **NAME**5604 `casinf` — complex arc sine functions5605 **SYNOPSIS**5606 `#include <complex.h>`5607 `float complex casinf(float complex z);`5608 **DESCRIPTION**5609 Refer to *casin()*.

5610 **NAME**

5611           casinh, casinhf, casinhl — complex arc hyperbolic sine functions

5612 **SYNOPSIS**

5613           #include &lt;complex.h&gt;

5614           double complex casinh(double complex z);

5615           float complex casinhf(float complex z);

5616           long double complex casinhl(long double complex z);

5617 **DESCRIPTION**

5618 **CX**       The functionality described on this reference page is aligned with the ISO C standard. Any  
5619           conflict between the requirements described here and the ISO C standard is unintentional. This  
5620           volume of IEEE Std 1003.1-200x defers to the ISO C standard.

5621           These functions shall compute the complex arc hyperbolic sine of  $z$ , with branch cuts outside the  
5622           interval  $[-i, +i]$  along the imaginary axis.

5623 **RETURN VALUE**

5624           These functions shall return the complex arc hyperbolic sine value, in the range of a strip  
5625           mathematically unbounded along the real axis and in the interval  $[-i\pi/2, +i\pi/2]$  along the  
5626           imaginary axis.

5627 **ERRORS**

5628           No errors are defined.

5629 **EXAMPLES**

5630           None.

5631 **APPLICATION USAGE**

5632           None.

5633 **RATIONALE**

5634           None.

5635 **FUTURE DIRECTIONS**

5636           None.

5637 **SEE ALSO**5638           `csinh()`, the Base Definitions volume of IEEE Std 1003.1-200x, <complex.h>5639 **CHANGE HISTORY**

5640           First released in Issue 6. Derived from the ISO/IEC 9899:1999 standard.

5641 **NAME**5642        **casinl** — complex arc sine functions5643 **SYNOPSIS**

5644        #include &lt;complex.h&gt;

5645        long double complex casinl(long double complex *z*);5646 **DESCRIPTION**5647        Refer to *casin()*.



5648 **NAME**

5649           catan, catanf, catanl — complex arc tangent functions

5650 **SYNOPSIS**

5651           #include &lt;complex.h&gt;

5652           double complex catan(double complex z);

5653           float complex catanf(float complex z);

5654           long double complex catanl(long double complex z);

5655 **DESCRIPTION**

5656 **CX**       The functionality described on this reference page is aligned with the ISO C standard. Any  
5657 conflict between the requirements described here and the ISO C standard is unintentional. This  
5658 volume of IEEE Std 1003.1-200x defers to the ISO C standard.

5659       These functions shall compute the complex arc tangent of  $z$ , with branch cuts outside the  
5660 interval  $[-i, +i]$  along the imaginary axis.

5661 **RETURN VALUE**

5662       These functions shall return the complex arc tangent value, in the range of a strip  
5663 mathematically unbounded along the imaginary axis and in the interval  $[-\pi/2, +\pi/2]$  along the  
5664 real axis.

5665 **ERRORS**

5666       No errors are defined.

5667 **EXAMPLES**

5668       None.

5669 **APPLICATION USAGE**

5670       None.

5671 **RATIONALE**

5672       None.

5673 **FUTURE DIRECTIONS**

5674       None.

5675 **SEE ALSO**

5676       ctan(), the Base Definitions volume of IEEE Std 1003.1-200x, &lt;complex.h&gt;

5677 **CHANGE HISTORY**

5678       First released in Issue 6. Derived from the ISO/IEC 9899:1999 standard.

5679 **NAME**5680        **catanf** — complex arc tangent functions5681 **SYNOPSIS**5682        `#include <complex.h>`5683        `float complex catanf(float complex z);`5684 **DESCRIPTION**5685        Refer to *catan()*.

5686 **NAME**

5687         catanh, catanhf, catanhl — complex arc hyperbolic tangent functions

5688 **SYNOPSIS**

5689         #include &lt;complex.h&gt;

5690         double complex catanh(double complex z);

5691         float complex catanhf(float complex z);

5692         long double complex catanhl(long double complex z);

5693 **DESCRIPTION**5694 **CX**         The functionality described on this reference page is aligned with the ISO C standard. Any  
5695 conflict between the requirements described here and the ISO C standard is unintentional. This  
5696 volume of IEEE Std 1003.1-200x defers to the ISO C standard.5697         These functions shall compute the complex arc hyperbolic tangent of  $z$ , with branch cuts outside  
5698 the interval  $[-1, +1]$  along the real axis.5699 **RETURN VALUE**5700         These functions shall return the complex arc hyperbolic tangent value, in the range of a strip  
5701 mathematically unbounded along the real axis and in the interval  $[-i\pi/2, +i\pi/2]$  along the  
5702 imaginary axis.5703 **ERRORS**

5704         No errors are defined.

5705 **EXAMPLES**

5706         None.

5707 **APPLICATION USAGE**

5708         None.

5709 **RATIONALE**

5710         None.

5711 **FUTURE DIRECTIONS**

5712         None.

5713 **SEE ALSO**5714         `ctanh()`, the Base Definitions volume of IEEE Std 1003.1-200x, <complex.h>5715 **CHANGE HISTORY**

5716         First released in Issue 6. Derived from the ISO/IEC 9899:1999 standard.

5717 **NAME**5718        **catanl** — complex arc tangent functions5719 **SYNOPSIS**5720        `#include <complex.h>`5721        `long double complex catanl(long double complex z);`5722 **DESCRIPTION**5723        Refer to *catan()*.

5724 **NAME**

5725           catclose — close a message catalog descriptor

5726 **SYNOPSIS**

5727 XSI       #include &lt;nl\_types.h&gt;

5728       int catclose(nl\_catd catd);

5729

5730 **DESCRIPTION**5731       The *catclose()* function shall close the message catalog identified by *catd*. If a file descriptor is  
5732       used to implement the type **nl\_catd**, that file descriptor shall be closed.5733 **RETURN VALUE**5734       Upon successful completion, *catclose()* shall return 0; otherwise, -1 shall be returned, and *errno*  
5735       set to indicate the error.5736 **ERRORS**5737       The *catclose()* function may fail if:

5738       [EBADF]           The catalog descriptor is not valid.

5739       [EINTR]           The *catclose()* function was interrupted by a signal.5740 **EXAMPLES**

5741       None.

5742 **APPLICATION USAGE**

5743       None.

5744 **RATIONALE**

5745       None.

5746 **FUTURE DIRECTIONS**

5747       None.

5748 **SEE ALSO**5749       *catgets()*, *catopen()*, the Base Definitions volume of IEEE Std 1003.1-200x, <nl\_types.h>5750 **CHANGE HISTORY**

5751       First released in Issue 2.

5752 **NAME**

5753       catgets — read a program message

5754 **SYNOPSIS**

5755 XSI       #include &lt;nl\_types.h&gt;

5756       char \*catgets(nl\_catd *catd*, int *set\_id*, int *msg\_id*, const char \**s*);

5757

5758 **DESCRIPTION**

5759       The *catgets()* function shall attempt to read message *msg\_id*, in set *set\_id*, from the message  
5760       catalog identified by *catd*. The *catd* argument is a message catalog descriptor returned from an  
5761       earlier call to *catopen()*. The *s* argument points to a default message string which shall be  
5762       returned by *catgets()* if it cannot retrieve the identified message.

5763       The *catgets()* function need not be reentrant. A function that is not required to be reentrant is not  
5764       required to be thread-safe.

5765 **RETURN VALUE**

5766       If the identified message is retrieved successfully, *catgets()* shall return a pointer to an internal  
5767       buffer area containing the null-terminated message string. If the call is unsuccessful for any  
5768       reason, *s* shall be returned and *errno* may be set to indicate the error.

5769 **ERRORS**5770       The *catgets()* function may fail if:5771       [EBADF]       The *catd* argument is not a valid message catalog descriptor open for reading.5772       [EBADMSG]     The message identified by *set\_id* and *msg\_id* in the specified message catalog  
5773       did not satisfy implementation-defined security criteria.5774       [EINTR]       The read operation was terminated due to the receipt of a signal, and no data  
5775       was transferred.5776       [EINVAL]      The message catalog identified by *catd* is corrupted.5777       [ENOMSG]     The message identified by *set\_id* and *msg\_id* is not in the message catalog.5778 **EXAMPLES**

5779       None.

5780 **APPLICATION USAGE**

5781       None.

5782 **RATIONALE**

5783       None.

5784 **FUTURE DIRECTIONS**

5785       None.

5786 **SEE ALSO**5787       *catclose()*, *catopen()*, the Base Definitions volume of IEEE Std 1003.1-200x, <nl\_types.h>5788 **CHANGE HISTORY**

5789       First released in Issue 2.

5790 **Issue 5**

5791       A note indicating that this function need not be reentrant is added to the DESCRIPTION.

5792 **Issue 6**

5793 In the DESCRIPTION, the note about reentrancy is expanded to cover thread-safety.

## 5794 NAME

5795 catopen — open a message catalog

## 5796 SYNOPSIS

5797 xSI #include &lt;nl\_types.h&gt;

5798 nl\_catd catopen(const char \*name, int oflag);

5799

## 5800 DESCRIPTION

5801 The *catopen()* function shall open a message catalog and return a message catalog descriptor.  
 5802 The *name* argument specifies the name of the message catalog to be opened. If *name* contains a  
 5803 `'/'`, then *name* specifies a complete name for the message catalog. Otherwise, the environment  
 5804 variable *NLSPATH* is used with *name* substituted for the `%N` conversion specification (see the  
 5805 Base Definitions volume of IEEE Std 1003.1-200x, Chapter 8, Environment Variables). If  
 5806 *NLSPATH* exists in the environment when the process starts, then if the process has appropriate  
 5807 privileges, the behavior of *catopen()* is undefined. If *NLSPATH* does not exist in the environment,  
 5808 or if a message catalog cannot be found in any of the components specified by *NLSPATH*, then  
 5809 an implementation-defined default path shall be used. This default may be affected by the  
 5810 setting of *LC\_MESSAGES* if the value of *oflag* is `NL_CAT_LOCALE`, or the *LANG* environment  
 5811 variable if *oflag* is 0.

5812 A message catalog descriptor shall remain valid in a process until that process closes it, or a  
 5813 successful call to one of the *exec* functions. A change in the setting of the *LC\_MESSAGES*  
 5814 category may invalidate existing open catalogs.

5815 If a file descriptor is used to implement message catalog descriptors, the `FD_CLOEXEC` flag  
 5816 shall be set; see <*fcntl.h*>.

5817 If the value of the *oflag* argument is 0, the *LANG* environment variable is used to locate the  
 5818 catalog without regard to the *LC\_MESSAGES* category. If the *oflag* argument is  
 5819 `NL_CAT_LOCALE`, the *LC\_MESSAGES* category is used to locate the message catalog (see the  
 5820 Base Definitions volume of IEEE Std 1003.1-200x, Section 8.2, Internationalization Variables).

## 5821 RETURN VALUE

5822 Upon successful completion, *catopen()* shall return a message catalog descriptor for use on  
 5823 subsequent calls to *catgets()* and *catclose()*. Otherwise, *catopen()* shall return `(nl_catd) -1` and set  
 5824 *errno* to indicate the error.

## 5825 ERRORS

5826 The *catopen()* function may fail if:

5827 [EACCES] Search permission is denied for the component of the path prefix of the  
 5828 message catalog or read permission is denied for the message catalog.

5829 [EMFILE] {OPEN\_MAX} file descriptors are currently open in the calling process.

5830 [ENAMETOOLONG]

5831 The length of a pathname of the message catalog exceeds {PATH\_MAX} or a  
 5832 pathname component is longer than {NAME\_MAX}.

5833 [ENAMETOOLONG]

5834 Pathname resolution of a symbolic link produced an intermediate result  
 5835 whose length exceeds {PATH\_MAX}.

5836 [ENFILE] Too many files are currently open in the system.

5837 [ENOENT] The message catalog does not exist or the *name* argument points to an empty  
 5838 string.



- 5839 [ENOMEM] Insufficient storage space is available.  
5840 [ENOTDIR] A component of the path prefix of the message catalog is not a directory.

**5841 EXAMPLES**

5842 None.

**5843 APPLICATION USAGE**

5844 Some implementations of *catopen()* use *malloc()* to allocate space for internal buffer areas. The  
5845 *catopen()* function may fail if there is insufficient storage space available to accommodate these  
5846 buffers.

5847 Conforming applications must assume that message catalog descriptors are not valid after a call  
5848 to one of the *exec* functions.

5849 Application writers should be aware that guidelines for the location of message catalogs have  
5850 not yet been developed. Therefore they should take care to avoid conflicting with catalogs used  
5851 by other applications and the standard utilities.

**5852 RATIONALE**

5853 None.

**5854 FUTURE DIRECTIONS**

5855 None.

**5856 SEE ALSO**

5857 *catclose()*, *catgets()*, the Base Definitions volume of IEEE Std 1003.1-200x, `<fcntl.h>`,  
5858 `<nl_types.h>`, the Shell and Utilities volume of IEEE Std 1003.1-200x

**5859 CHANGE HISTORY**

5860 First released in Issue 2.

5861 **NAME**

5862           cbrt, cbrtf, cbrtl — cube root functions

5863 **SYNOPSIS**

5864           #include &lt;math.h&gt;

5865           double cbrt(double x);

5866           float cbrtf(float x);

5867           long double cbrtl(long double x);

5868 **DESCRIPTION**

5869 **CX**       The functionality described on this reference page is aligned with the ISO C standard. Any  
5870 conflict between the requirements described here and the ISO C standard is unintentional. This  
5871 volume of IEEE Std 1003.1-200x defers to the ISO C standard.

5872       These functions shall compute the real cube root of their argument *x*.5873 **RETURN VALUE**5874       Upon successful completion, these functions shall return the cube root of *x*.5875 **MX**       If *x* is NaN, a NaN shall be returned.5876       If *x* is  $\pm 0$ , or  $\pm\text{Inf}$ , *x* shall be returned.5877 **ERRORS**

5878       No errors are defined.

5879 **EXAMPLES**

5880       None.

5881 **APPLICATION USAGE**

5882       None.

5883 **RATIONALE**

5884       For some applications, a true cube root function, which returns negative results for negative  
5885 arguments, is more appropriate than *pow(x, 1.0/3.0)*, which returns a NaN for *x* less than 0.

5886 **FUTURE DIRECTIONS**

5887       None.

5888 **SEE ALSO**

5889       The Base Definitions volume of IEEE Std 1003.1-200x, &lt;math.h&gt;

5890 **CHANGE HISTORY**

5891       First released in Issue 4, Version 2.

5892 **Issue 5**

5893       Moved from X/OPEN UNIX extension to BASE.

5894 **Issue 6**5895       The *cbrt()* function is no longer marked as an extension.5896       The *cbrtf()* and *cbrtl()* functions are added for alignment with the ISO/IEC 9899:1999 standard.

5897       The DESCRIPTION, RETURN VALUE, ERRORS, and APPLICATION USAGE sections are  
5898 revised to align with the ISO/IEC 9899:1999 standard.

5899       IEC 60559:1989 standard floating-point extensions over the ISO/IEC 9899:1999 standard are  
5900 marked.

5901 **NAME**

5902 ccos, ccosf, ccosl — complex cosine functions

5903 **SYNOPSIS**

5904 #include &lt;complex.h&gt;

5905 double complex ccos(double complex z);

5906 float complex ccosf(float complex z);

5907 long double complex ccosl(long double complex z);

5908 **DESCRIPTION**

5909 cx The functionality described on this reference page is aligned with the ISO C standard. Any  
5910 conflict between the requirements described here and the ISO C standard is unintentional. This  
5911 volume of IEEE Std 1003.1-200x defers to the ISO C standard.

5912 These functions shall compute the complex cosine of  $z$ .5913 **RETURN VALUE**

5914 These functions shall return the complex cosine value.

5915 **ERRORS**

5916 No errors are defined.

5917 **EXAMPLES**

5918 None.

5919 **APPLICATION USAGE**

5920 None.

5921 **RATIONALE**

5922 None.

5923 **FUTURE DIRECTIONS**

5924 None.

5925 **SEE ALSO**5926 `cacos()`, the Base Definitions volume of IEEE Std 1003.1-200x, <complex.h>5927 **CHANGE HISTORY**

5928 First released in Issue 6. Derived from the ISO/IEC 9899:1999 standard.

5929 **NAME**5930 `ccosf` — complex cosine functions5931 **SYNOPSIS**5932 `#include <complex.h>`5933 `float complex ccosf(float complex z);`5934 **DESCRIPTION**5935 Refer to `ccos()`.

5936 **NAME**

5937 ccosh, ccoshf, ccoshl — complex hyperbolic cosine functions

5938 **SYNOPSIS**

5939 #include &lt;complex.h&gt;

5940 double complex ccosh(double complex z);

5941 float complex ccoshf(float complex z);

5942 long double complex ccoshl(long double complex z);

5943 **DESCRIPTION**

5944 cx The functionality described on this reference page is aligned with the ISO C standard. Any  
5945 conflict between the requirements described here and the ISO C standard is unintentional. This  
5946 volume of IEEE Std 1003.1-200x defers to the ISO C standard.

5947 These functions shall compute the complex hyperbolic cosine of  $z$ .5948 **RETURN VALUE**

5949 These functions shall return the complex hyperbolic cosine value.

5950 **ERRORS**

5951 No errors are defined.

5952 **EXAMPLES**

5953 None.

5954 **APPLICATION USAGE**

5955 None.

5956 **RATIONALE**

5957 None.

5958 **FUTURE DIRECTIONS**

5959 None.

5960 **SEE ALSO**5961 *cacosh()*, the Base Definitions volume of IEEE Std 1003.1-200x, <complex.h>5962 **CHANGE HISTORY**

5963 First released in Issue 6. Derived from the ISO/IEC 9899:1999 standard.

5964 **NAME**

5965       ccosl — complex cosine functions

5966 **SYNOPSIS**

5967       #include &lt;complex.h&gt;

5968       long double complex ccosl(long double complex z);

5969 **DESCRIPTION**5970       Refer to *ccos()*.

5971 **NAME**

5972       ceil, ceilf, ceill — ceiling value function

5973 **SYNOPSIS**

5974       #include &lt;math.h&gt;

5975       double ceil(double x);

5976       float ceilf(float x);

5977       long double ceill(long double x);

5978 **DESCRIPTION**

5979 **CX**       The functionality described on this reference page is aligned with the ISO C standard. Any  
 5980 conflict between the requirements described here and the ISO C standard is unintentional. This  
 5981 volume of IEEE Std 1003.1-200x defers to the ISO C standard.

5982       These functions shall compute the smallest integral value not less than *x*.

5983       An application wishing to check for error situations should set *errno* to zero and call  
 5984 *feclearexcept*(FE\_ALL\_EXCEPT) before calling these functions. On return, if *errno* is non-zero or  
 5985 *fetestexcept*(FE\_INVALID | FE\_DIVBYZERO | FE\_OVERFLOW | FE\_UNDERFLOW) is non-  
 5986 zero, an error has occurred.

5987 **RETURN VALUE**

5988       Upon successful completion, *ceil()*, *ceilf()*, and *ceill()* shall return the smallest integral value not  
 5989 less than *x*, expressed as a type **double**, **float**, or **long double**, respectively.

5990 **MX**       If *x* is NaN, a NaN shall be returned.5991       If *x* is  $\pm 0$ , or  $\pm \text{Inf}$ , *x* shall be returned.

5992 **XSI**       If the correct value would cause overflow, a range error shall occur and *ceil()*, *ceilf()*, and *ceill()*  
 5993 shall return the value of the macro HUGE\_VAL, HUGE\_VALF, and HUGE\_VALL, respectively.

5994 **ERRORS**

5995       These functions shall fail if:

5996 **XSI**       **Range Error**       The result overflows.

5997       If the integer expression (math\_errhandling & MATH\_ERRNO) is non-zero, |  
 5998 then *errno* shall be set to [ERANGE]. If the integer expression |  
 5999 (math\_errhandling & MATH\_ERREXCEPT) is non-zero, then the overflow |  
 6000 floating-point exception shall be raised. |

6001 **EXAMPLES**

6002       None.

6003 **APPLICATION USAGE**

6004       The integral value returned by these functions need not be expressible as an **int** or **long**. The  
 6005 return value should be tested before assigning it to an integer type to avoid the undefined results  
 6006 of an integer overflow.

6007       The *ceil()* function can only overflow when the floating-point representation has  
 6008 DBL\_MANT\_DIG > DBL\_MAX\_EXP.

6009       On error, the expressions (math\_errhandling & MATH\_ERRNO) and (math\_errhandling &  
 6010 MATH\_ERREXCEPT) are independent of each other, but at least one of them must be non-zero.

6011 **RATIONALE**

6012       None.

6013 **FUTURE DIRECTIONS**

6014       None.

6015 **SEE ALSO**6016       *feclearexcept()*, *fetestexcept()*, *floor()*, *isnan()*, the Base Definitions volume of IEEE Std 1003.1-200x, |

6017       Section 4.18, Treatment of Error Conditions for Mathematical Functions, &lt;math.h&gt; |

6018 **CHANGE HISTORY**

6019       First released in Issue 1. Derived from Issue 1 of the SVID.

6020 **Issue 5**6021       The DESCRIPTION is updated to indicate how an application should check for an error. This  
6022       text was previously published in the APPLICATION USAGE section.6023 **Issue 6**6024       The *ceilf()* and *ceilll()* functions are added for alignment with the ISO/IEC 9899:1999 standard.6025       The DESCRIPTION, RETURN VALUE, ERRORS, and APPLICATION USAGE sections are  
6026       revised to align with the ISO/IEC 9899:1999 standard.6027       IEC 60559:1989 standard floating-point extensions over the ISO/IEC 9899:1999 standard are  
6028       marked.



6029 **NAME**

6030 cexp, cexpf, cexpl — complex exponential functions

6031 **SYNOPSIS**

6032 #include &lt;complex.h&gt;

6033 double complex cexp(double complex z);

6034 float complex cexpf(float complex z);

6035 long double complex cexpl(long double complex z);

6036 **DESCRIPTION**6037 cx The functionality described on this reference page is aligned with the ISO C standard. Any  
6038 conflict between the requirements described here and the ISO C standard is unintentional. This  
6039 volume of IEEE Std 1003.1-200x defers to the ISO C standard.6040 These functions shall compute the complex exponent of  $z$ , defined as  $e^z$ .6041 **RETURN VALUE**6042 These functions shall return the complex exponential value of  $z$ .6043 **ERRORS**

6044 No errors are defined.

6045 **EXAMPLES**

6046 None.

6047 **APPLICATION USAGE**

6048 None.

6049 **RATIONALE**

6050 None.

6051 **FUTURE DIRECTIONS**

6052 None.

6053 **SEE ALSO**6054 *clog()*, the Base Definitions volume of IEEE Std 1003.1-200x, <complex.h>6055 **CHANGE HISTORY**

6056 First released in Issue 6. Derived from the ISO/IEC 9899:1999 standard.

6057 **NAME**

6058 cfgetispeed — get input baud rate

6059 **SYNOPSIS**

6060 #include &lt;termios.h&gt;

6061 speed\_t cfgetispeed(const struct termios \*termios\_p);

6062 **DESCRIPTION**6063 The *cfgetispeed()* function shall extract the input baud rate from the **termios** structure to which  
6064 the *termios\_p* argument points.6065 This function shall return exactly the value in the **termios** data structure, without interpretation.6066 **RETURN VALUE**6067 Upon successful completion, *cfgetispeed()* shall return a value of type **speed\_t** representing the  
6068 input baud rate.6069 **ERRORS**

6070 No errors are defined.

6071 **EXAMPLES**

6072 None.

6073 **APPLICATION USAGE**

6074 None.

6075 **RATIONALE**6076 The term *baud* is used historically here, but is not technically correct. This is properly “bits per  
6077 second”, which may not be the same as baud. However, the term is used because of the  
6078 historical usage and understanding.6079 The *cfgetospeed()*, *cfgetispeed()*, *cfsetospeed()*, and *cfsetispeed()* functions do not take arguments as  
6080 numbers, but rather as symbolic names. There are two reasons for this:

- 6081 1. Historically, numbers were not used because of the way the rate was stored in the data
- 
- 6082 structure. This is retained even though a function is now used.
- 
- 6083 2. More importantly, only a limited set of possible rates is at all portable, and this constrains
- 
- 6084 the application to that set.

6085 There is nothing to prevent an implementation to accept, as an extension, a number (such as 126)  
6086 if it wished, and because the encoding of the Bxxx symbols is not specified, this can be done so  
6087 no ambiguity is introduced.6088 Setting the input baud rate to zero was a mechanism to allow for split baud rates. Clarifications  
6089 in this volume of IEEE Std 1003.1-200x have made it possible to determine whether split rates are  
6090 supported and to support them without having to treat zero as a special case. Since this  
6091 functionality is also confusing, it has been declared obsolescent. The 0 argument referred to is  
6092 the literal constant 0, not the symbolic constant B0. This volume of IEEE Std 1003.1-200x does  
6093 not preclude B0 from being defined as the value 0; in fact, implementations would likely benefit  
6094 from the two being equivalent. This volume of IEEE Std 1003.1-200x does not fully specify  
6095 whether the previous *cfsetispeed()* value is retained after a *tcgetattr()* as the actual value or as  
6096 zero. Therefore, conforming applications should always set both the input speed and output  
6097 speed when setting either.6098 In historical implementations, the baud rate information is traditionally kept in **c\_cflag**.  
6099 Applications should be written to presume that this might be the case (and thus not blindly copy  
6100 **c\_cflag**), but not to rely on it in case it is in some other field of the structure. Setting the **c\_cflag**  
6101 field absolutely after setting a baud rate is a non-portable action because of this. In general, the

6102 unused parts of the flag fields might be used by the implementation and should not be blindly  
6103 copied from the descriptions of one terminal device to another.

6104 **FUTURE DIRECTIONS**

6105 None.

6106 **SEE ALSO**

6107 *cfgetospeed()*, *cfsetispeed()*, *cfsetospeed()*, *tcgetattr()*, the Base Definitions volume of  
6108 IEEE Std 1003.1-200x, <**termios.h**>, the Base Definitions volume of IEEE Std 1003.1-200x,  
6109 Chapter 11, General Terminal Interface

6110 **CHANGE HISTORY**

6111 First released in Issue 3.

6112 Entry included for alignment with the POSIX.1-1988 standard.

6113 **NAME**

6114 cfgetospeed — get output baud rate

6115 **SYNOPSIS**

6116 #include <termios.h>

6117 speed\_t cfgetospeed(const struct termios \*termios\_p);

6118 **DESCRIPTION**

6119 The *cfgetospeed()* function shall extract the output baud rate from the **termios** structure to which  
6120 the *termios\_p* argument points.

6121 This function shall return exactly the value in the **termios** data structure, without interpretation.

6122 **RETURN VALUE**

6123 Upon successful completion, *cfgetospeed()* shall return a value of type **speed\_t** representing the  
6124 output baud rate.

6125 **ERRORS**

6126 No errors are defined.

6127 **EXAMPLES**

6128 None.

6129 **APPLICATION USAGE**

6130 None.

6131 **RATIONALE**

6132 Refer to *cfgetispeed()*.

6133 **FUTURE DIRECTIONS**

6134 None.

6135 **SEE ALSO**

6136 *cfgetispeed()*, *cfsetispeed()*, *cfsetospeed()*, *tcgetattr()*, the Base Definitions volume of  
6137 IEEE Std 1003.1-200x, <**termios.h**>, the Base Definitions volume of IEEE Std 1003.1-200x,  
6138 Chapter 11, General Terminal Interface

6139 **CHANGE HISTORY**

6140 First released in Issue 3.

6141 Entry included for alignment with the POSIX.1-1988 standard.

6142 **NAME**

6143 cfsetispeed — set input baud rate

6144 **SYNOPSIS**

6145 #include &lt;termios.h&gt;

6146 int cfsetispeed(struct termios \*termios\_p, speed\_t speed);

6147 **DESCRIPTION**6148 The *cfsetispeed()* function shall set the input baud rate stored in the structure pointed to by  
6149 *termios\_p* to *speed*.6150 There shall be no effect on the baud rates set in the hardware until a subsequent successful call  
6151 to *tcsetattr()* with the same **termios** structure. Similarly, errors resulting from attempts to set  
6152 baud rates not supported by the terminal device need not be detected until the *tcsetattr()*  
6153 function is called.6154 **RETURN VALUE**6155 Upon successful completion, *cfsetispeed()* shall return 0; otherwise, -1 shall be returned, and  
6156 *errno* may be set to indicate the error.6157 **ERRORS**6158 The *cfsetispeed()* function may fail if:6159 [EINVAL] The *speed* value is not a valid baud rate.6160 [EINVAL] The value of *speed* is outside the range of possible speed values as specified in  
6161 <termios.h>.6162 **EXAMPLES**

6163 None.

6164 **APPLICATION USAGE**

6165 None.

6166 **RATIONALE**6167 Refer to *cfgetispeed()*.6168 **FUTURE DIRECTIONS**

6169 None.

6170 **SEE ALSO**6171 *cfgetispeed()*, *cfgetospeed()*, *cfsetospeed()*, *tcsetattr()*, the Base Definitions volume of  
6172 IEEE Std 1003.1-200x, <termios.h>, the Base Definitions volume of IEEE Std 1003.1-200x,  
6173 Chapter 11, General Terminal Interface6174 **CHANGE HISTORY**

6175 First released in Issue 3.

6176 Entry included for alignment with the POSIX.1-1988 standard.

6177 **Issue 6**6178 The following new requirements on POSIX implementations derive from alignment with the  
6179 Single UNIX Specification:

- 6180
- The optional setting of *errno* and the [EINVAL] error conditions are added.

6181 **NAME**

6182 cfsetospeed — set output baud rate

6183 **SYNOPSIS**

6184 #include &lt;termios.h&gt;

6185 int cfsetospeed(struct termios \*termios\_p, speed\_t speed);

6186 **DESCRIPTION**6187 The *cfsetospeed()* function shall set the output baud rate stored in the structure pointed to by  
6188 *termios\_p* to *speed*.6189 There shall be no effect on the baud rates set in the hardware until a subsequent successful call  
6190 to *tcsetattr()* with the same **termios** structure. Similarly, errors resulting from attempts to set  
6191 baud rates not supported by the terminal device need not be detected until the *tcsetattr()*  
6192 function is called.6193 **RETURN VALUE**6194 Upon successful completion, *cfsetospeed()* shall return 0; otherwise, it shall return -1 and *errno*  
6195 may be set to indicate the error.6196 **ERRORS**6197 The *cfsetospeed()* function may fail if:6198 [EINVAL] The *speed* value is not a valid baud rate.6199 [EINVAL] The value of *speed* is outside the range of possible speed values as specified in  
6200 <**termios.h**>.6201 **EXAMPLES**

6202 None.

6203 **APPLICATION USAGE**

6204 None.

6205 **RATIONALE**6206 Refer to *cfgetispeed()*.6207 **FUTURE DIRECTIONS**

6208 None.

6209 **SEE ALSO**6210 *cfgetispeed()*, *cfgetospeed()*, *cfsetispeed()*, *tcsetattr()*, the Base Definitions volume of  
6211 IEEE Std 1003.1-200x, <**termios.h**>, the Base Definitions volume of IEEE Std 1003.1-200x,  
6212 Chapter 11, General Terminal Interface6213 **CHANGE HISTORY**

6214 First released in Issue 3.

6215 Entry included for alignment with the POSIX.1-1988 standard.

6216 **Issue 6**6217 The following new requirements on POSIX implementations derive from alignment with the  
6218 Single UNIX Specification:

- 6219
- The optional setting of *errno* and the [EINVAL] error conditions are added.

6220 **NAME**

6221           chdir — change working directory

6222 **SYNOPSIS**

6223           #include &lt;unistd.h&gt;

6224           int chdir(const char \*path);

6225 **DESCRIPTION**

6226           The *chdir()* function shall cause the directory named by the pathname pointed to by the *path* |  
 6227           argument to become the current working directory; that is, the starting point for path searches |  
 6228           for pathnames not beginning with '/'. |

6229 **RETURN VALUE**

6230           Upon successful completion, 0 shall be returned. Otherwise, -1 shall be returned, the current |  
 6231           working directory shall remain unchanged, and *errno* shall be set to indicate the error.

6232 **ERRORS**6233           The *chdir()* function shall fail if:

6234           [EACCES]           Search permission is denied for any component of the pathname. |

6235           [ELOOP]           A loop exists in symbolic links encountered during resolution of the *path* |  
6236           argument.

6237           [ENAMETOOLONG]

6238                               The length of the *path* argument exceeds {PATH\_MAX} or a pathname |  
6239                               component is longer than {NAME\_MAX}.6240           [ENOENT]           A component of *path* does not name an existing directory or *path* is an empty |  
6241           string.

6242           [ENOTDIR]          A component of the pathname is not a directory. |

6243           The *chdir()* function may fail if:6244           [ELOOP]           More than {SYMLOOP\_MAX} symbolic links were encountered during |  
6245           resolution of the *path* argument.

6246           [ENAMETOOLONG]

6247                               As a result of encountering a symbolic link in resolution of the *path* argument, |  
6248                               the length of the substituted pathname string exceeded {PATH\_MAX}.6249 **EXAMPLES**6250           **Changing the Current Working Directory**6251           The following example makes the value pointed to by **directory**, **/tmp**, the current working |  
6252           directory.

6253           #include &lt;unistd.h&gt;

6254           ...

6255           char \*directory = "/tmp";

6256           int ret;

6257           ret = chdir (directory);

6258 **APPLICATION USAGE**

6259 None.

6260 **RATIONALE**6261 The *chdir()* function only affects the working directory of the current process.6262 **FUTURE DIRECTIONS**

6263 None.

6264 **SEE ALSO**6265 *getcwd()*, the Base Definitions volume of IEEE Std 1003.1-200x, <**unistd.h**>6266 **CHANGE HISTORY**

6267 First released in Issue 1. Derived from Issue 1 of the SVID.

6268 **Issue 6**

6269 The APPLICATION USAGE section is added.

6270 The following new requirements on POSIX implementations derive from alignment with the |  
6271 Single UNIX Specification:

- 6272
- The [ELOOP] mandatory error condition is added.
  - A second [ENAMETOOLONG] is added as an optional error condition.

6274 The following changes were made to align with the IEEE P1003.1a draft standard:

- 6275
- The [ELOOP] optional error condition is added.



6276 **NAME**

6277 chmod — change mode of a file

6278 **SYNOPSIS**

6279 #include &lt;sys/stat.h&gt;

6280 int chmod(const char \*path, mode\_t mode);

6281 **DESCRIPTION**

6282 XSI The *chmod()* function shall change S\_ISUID, S\_ISGID, S\_ISVTX, and the file permission bits of |  
 6283 the file named by the pathname pointed to by the *path* argument to the corresponding bits in the |  
 6284 *mode* argument. The application shall ensure that the effective user ID of the process matches the |  
 6285 owner of the file or the process has appropriate privileges in order to do this.

6286 XSI S\_ISUID, S\_ISGID, S\_ISVTX, and the file permission bits are described in &lt;sys/stat.h&gt;.

6287 If the calling process does not have appropriate privileges, and if the group ID of the file does |  
 6288 not match the effective group ID or one of the supplementary group IDs and if the file is a |  
 6289 regular file, bit S\_ISGID (set-group-ID on execution) in the file's mode shall be cleared upon |  
 6290 successful return from *chmod()*.

6291 Additional implementation-defined restrictions may cause the S\_ISUID and S\_ISGID bits in |  
 6292 *mode* to be ignored.

6293 The effect on file descriptors for files open at the time of a call to *chmod()* is implementation- |  
 6294 defined.

6295 Upon successful completion, *chmod()* shall mark for update the *st\_ctime* field of the file.6296 **RETURN VALUE**

6297 Upon successful completion, 0 shall be returned; otherwise, -1 shall be returned and *errno* set to |  
 6298 indicate the error. If -1 is returned, no change to the file mode occurs.

6299 **ERRORS**6300 The *chmod()* function shall fail if:

6301 [EACCES] Search permission is denied on a component of the path prefix.

6302 [ELOOP] A loop exists in symbolic links encountered during resolution of the *path* |  
 6303 argument.

6304 [ENAMETOOLONG]

6305 The length of the *path* argument exceeds {PATH\_MAX} or a pathname |  
 6306 component is longer than {NAME\_MAX}.

6307 [ENOTDIR] A component of the path prefix is not a directory.

6308 [ENOENT] A component of *path* does not name an existing file or *path* is an empty string.6309 [EPERM] The effective user ID does not match the owner of the file and the process |  
 6310 does not have appropriate privileges.

6311 [EROFS] The named file resides on a read-only file system.

6312 The *chmod()* function may fail if:

6313 [EINTR] A signal was caught during execution of the function.

6314 [EINVAL] The value of the *mode* argument is invalid.6315 [ELOOP] More than {SYMLOOP\_MAX} symbolic links were encountered during |  
 6316 resolution of the *path* argument.

6317 [ENAMETOOLONG]  
6318 As a result of encountering a symbolic link in resolution of the *path* argument, |  
6319 the length of the substituted pathname strings exceeded {PATH\_MAX}. |

## 6320 EXAMPLES

### 6321 **Setting Read Permissions for User, Group, and Others**

6322 The following example sets read permissions for the owner, group, and others.

```
6323 #include <sys/stat.h>
6324 const char *path;
6325 ...
6326 chmod(path, S_IRUSR|S_IRGRP|S_IROTH);
```

### 6327 **Setting Read, Write, and Execute Permissions for the Owner Only**

6328 The following example sets read, write, and execute permissions for the owner, and no  
6329 permissions for group and others.

```
6330 #include <sys/stat.h>
6331 const char *path;
6332 ...
6333 chmod(path, S_IRWXU);
```

### 6334 **Setting Different Permissions for Owner, Group, and Other**

6335 The following example sets owner permissions for CHANGEFILE to read, write, and execute,  
6336 group permissions to read and execute, and other permissions to read.

```
6337 #include <sys/stat.h>
6338 #define CHANGEFILE "/etc/myfile"
6339 ...
6340 chmod(CHANGEFILE, S_IRWXU|S_IRGRP|S_IXGRP|S_IROTH);
```

### 6341 **Setting and Checking File Permissions**

6342 The following example sets the file permission bits for a file named */home/cnd/mod1*, then calls  
6343 the *stat()* function to verify the permissions.

```
6344 #include <sys/types.h>
6345 #include <sys/stat.h>
6346 int status;
6347 struct stat buffer
6348 ...
6349 chmod("home/cnd/mod1", S_IRWXU|S_IRWXG|S_IROTH|S_IWOTH);
6350 status = stat("home/cnd/mod1", &buffer);
```

## 6351 APPLICATION USAGE

6352 In order to ensure that the *S\_ISUID* and *S\_ISGID* bits are set, an application requiring this should  
6353 use *stat()* after a successful *chmod()* to verify this.

6354 Any file descriptors currently open by any process on the file could possibly become invalid if  
6355 the mode of the file is changed to a value which would deny access to that process. One

6356 situation where this could occur is on a stateless file system. This behavior will not occur in a  
6357 conforming environment.

6358 **RATIONALE**

6359 This volume of IEEE Std 1003.1-200x specifies that the S\_ISGID bit is cleared by *chmod()* on a  
6360 regular file under certain conditions. This is specified on the assumption that regular files may  
6361 be executed, and the system should prevent users from making executable *setgid()* files perform  
6362 with privileges that the caller does not have. On implementations that support execution of  
6363 other file types, the S\_ISGID bit should be cleared for those file types under the same  
6364 circumstances.

6365 Implementations that use the S\_ISUID bit to indicate some other function (for example,  
6366 mandatory record locking) on non-executable files need not clear this bit on writing. They  
6367 should clear the bit for executable files and any other cases where the bit grants special powers  
6368 to processes that change the file contents. Similar comments apply to the S\_ISGID bit.

6369 **FUTURE DIRECTIONS**

6370 None.

6371 **SEE ALSO**

6372 *chown()*, *mkdir()*, *mkfifo()*, *open()*, *stat()*, *statvfs()*, the Base Definitions volume of  
6373 IEEE Std 1003.1-200x, `<sys/stat.h>`, `<sys/types.h>`

6374 **CHANGE HISTORY**

6375 First released in Issue 1. Derived from Issue 1 of the SVID.

6376 **Issue 6**

6377 The following new requirements on POSIX implementations derive from alignment with the |  
6378 Single UNIX Specification:

6379 • The requirement to include `<sys/types.h>` has been removed. Although `<sys/types.h>` was  
6380 required for conforming implementations of previous POSIX specifications, it was not  
6381 required for UNIX applications.

6382 • The [EINVAL] and [EINTR] optional error conditions are added.

6383 • A second [ENAMETOOLONG] is added as an optional error condition.

6384 The following changes were made to align with the IEEE P1003.1a draft standard:

6385 • The [ELOOP] optional error condition is added.

6386 The DESCRIPTION is updated to avoid use of the term “must” for application requirements.

6387 **NAME**

6388           chown — change owner and group of a file

6389 **SYNOPSIS**

6390           #include &lt;unistd.h&gt;

6391           int chown(const char \*path, uid\_t owner, gid\_t group);

6392 **DESCRIPTION**6393           The *chown()* function shall change the user and group ownership of a file. |6394           The *path* argument points to a pathname naming a file. The user ID and group ID of the named |  
6395           file shall be set to the numeric values contained in *owner* and *group*, respectively. |6396           Only processes with an effective user ID equal to the user ID of the file or with appropriate |  
6397           privileges may change the ownership of a file. If `_POSIX_CHOWN_RESTRICTED` is in effect for |  
6398           *path*:

- 6399
- 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 ID of the file, but without appropriate privileges, if and only if *owner* is equal to the file's user ID or `(uid_t)-1` and *group* is equal either to the calling process' effective group ID or to one of its supplementary group IDs.
- 6400
- 
- 6401
- 
- 6402
- 
- 6403

6404           If the specified file is a regular file, one or more of the `S_IXUSR`, `S_IXGRP`, or `S_IXOTH` bits of |  
6405           the file mode are set, and the process does not have appropriate privileges, the set-user-ID |  
6406           (`S_ISUID`) and set-group-ID (`S_ISGID`) bits of the file mode shall be cleared upon successful |  
6407           return from *chown()*. If the specified file is a regular file, one or more of the `S_IXUSR`, `S_IXGRP`, |  
6408           or `S_IXOTH` bits of the file mode are set, and the process has appropriate privileges, it is |  
6409           implementation-defined whether the set-user-ID and set-group-ID bits are altered. If the *chown()* |  
6410           function is successfully invoked on a file that is not a regular file and one or more of the |  
6411           `S_IXUSR`, `S_IXGRP`, or `S_IXOTH` bits of the file mode are set, the set-user-ID and set-group-ID |  
6412           bits may be cleared.6413           If *owner* or *group* is specified as `(uid_t)-1` or `(gid_t)-1`, respectively, the corresponding ID of the |  
6414           file shall not be changed. If both *owner* and *group* are `-1`, the times need not be updated. |6415           Upon successful completion, *chown()* shall mark for update the *st\_ctime* field of the file.6416 **RETURN VALUE**6417           Upon successful completion, 0 shall be returned; otherwise, `-1` shall be returned and *errno* set to |  
6418           indicate the error. If `-1` is returned, no changes are made in the user ID and group ID of the file.6419 **ERRORS**6420           The *chown()* function shall fail if:

- 6421           [EACCES]           Search permission is denied on a component of the path prefix.
- 
- 6422           [ELOOP]           A loop exists in symbolic links encountered during resolution of the
- path*
- |
- 
- 6423           argument.
- 
- 6424           [ENAMETOOLONG]       The length of the
- path*
- argument exceeds {PATH\_MAX} or a pathname |
- 
- 6425           component is longer than {NAME\_MAX}. |
- 
- 6426
- 
- 6427           [ENOTDIR]           A component of the path prefix is not a directory.
- 
- 6428           [ENOENT]           A component of
- path*
- does not name an existing file or
- path*
- is an empty string.

6429 [EPERM] The effective user ID does not match the owner of the file, or the calling  
 6430 process does not have appropriate privileges and  
 6431 `_POSIX_CHOWN_RESTRICTED` indicates that such privilege is required.

6432 [EROFS] The named file resides on a read-only file system.

6433 The `chown()` function may fail if:

6434 [EIO] An I/O error occurred while reading or writing to the file system.

6435 [EINTR] The `chown()` function was interrupted by a signal which was caught.

6436 [EINVAL] The owner or group ID supplied is not a value supported by the  
 6437 implementation.

6438 [ELOOP] More than `{SYMLOOP_MAX}` symbolic links were encountered during  
 6439 resolution of the *path* argument.

6440 [ENAMETOOLONG]

6441 As a result of encountering a symbolic link in resolution of the *path* argument, |  
 6442 the length of the substituted pathname string exceeded `{PATH_MAX}`. |

#### 6443 EXAMPLES

6444 None.

#### 6445 APPLICATION USAGE

6446 Although `chown()` can be used on some implementations by the file owner to change the owner |  
 6447 and group to any desired values, the only portable use of this function is to change the group of |  
 6448 a file to the effective GID of the calling process or to a member of its group set. |

#### 6449 RATIONALE

6450 System III and System V allow a user to give away files; that is, the owner of a file may change  
 6451 its user ID to anything. This is a serious problem for implementations that are intended to meet  
 6452 government security regulations. Version 7 and 4.3 BSD permit only the superuser to change the  
 6453 user ID of a file. Some government agencies (usually not ones concerned directly with security)  
 6454 find this limitation too confining. This volume of IEEE Std 1003.1-200x uses *may* to permit secure  
 6455 implementations while not disallowing System V.

6456 System III and System V allow the owner of a file to change the group ID to anything. Version 7  
 6457 permits only the superuser to change the group ID of a file. 4.3 BSD permits the owner to  
 6458 change the group ID of a file to its effective group ID or to any of the groups in the list of  
 6459 supplementary group IDs, but to no others.

6460 The POSIX.1-1990 standard requires that the `chown()` function invoked by a non-appropriate  
 6461 privileged process clear the `S_ISGID` and the `S_ISUID` bits for regular files, and permits them to  
 6462 be cleared for other types of files. This is so that changes in accessibility do not accidentally  
 6463 cause files to become security holes. Unfortunately, requiring these bits to be cleared on non-  
 6464 executable data files also clears the mandatory file locking bit (shared with `S_ISGID`), which is  
 6465 an extension on many implementations (it first appeared in System V). These bits should only be  
 6466 required to be cleared on regular files that have one or more of their execute bits set.

#### 6467 FUTURE DIRECTIONS

6468 None.

#### 6469 SEE ALSO

6470 `chmod()`, `pathconf()`, the Base Definitions volume of IEEE Std 1003.1-200x, `<sys/types.h>`,  
 6471 `<unistd.h>`

6472 **CHANGE HISTORY**

6473 First released in Issue 1. Derived from Issue 1 of the SVID.

6474 **Issue 6**

6475 The following changes are made for alignment with the ISO POSIX-1: 1996 standard:

- 6476 • The wording describing the optional dependency on `_POSIX_CHOWN_RESTRICTED` is  
6477 restored.
- 6478 • The [EPERM] error is restored as an error dependent on `_POSIX_CHOWN_RESTRICTED`. |  
6479 This is since its operand is a pathname and applications should be aware that the error may |  
6480 not occur for that pathname if the file system does not support |  
6481 `_POSIX_CHOWN_RESTRICTED`. |

6482 The following new requirements on POSIX implementations derive from alignment with the  
6483 Single UNIX Specification:

- 6484 • The requirement to include `<sys/types.h>` has been removed. Although `<sys/types.h>` was  
6485 required for conforming implementations of previous POSIX specifications, it was not  
6486 required for UNIX applications.
- 6487 • The value for *owner* of `(uid_t)-1` allows the use of `-1` by the owner of a file to change the  
6488 group ID only. A corresponding change is made for group. |
- 6489 • The [ELOOP] mandatory error condition is added.
- 6490 • The [EIO] and [EINTR] optional error conditions are added.
- 6491 • A second [ENAMETOOLONG] is added as an optional error condition.

6492 The following changes were made to align with the IEEE P1003.1a draft standard:

- 6493 • Clarification is added that the `S_ISUID` and `S_ISGID` bits do not need to be cleared when the  
6494 process has appropriate privileges.
- 6495 • The [ELOOP] optional error condition is added.

6496 **NAME**6497 `cimag`, `cimagf`, `cimagl` — complex imaginary functions6498 **SYNOPSIS**6499 `#include <complex.h>`6500 `double cimag(double complex z);`6501 `float cimagf(float complex z);`6502 `long double cimagl(long double complex z);`6503 **DESCRIPTION**

6504 **CX** The functionality described on this reference page is aligned with the ISO C standard. Any  
6505 conflict between the requirements described here and the ISO C standard is unintentional. This  
6506 volume of IEEE Std 1003.1-200x defers to the ISO C standard.

6507 These functions shall compute the imaginary part of  $z$ .6508 **RETURN VALUE**

6509 These functions shall return the imaginary part value (as a real).

6510 **ERRORS**

6511 No errors are defined.

6512 **EXAMPLES**

6513 None.

6514 **APPLICATION USAGE**6515 For a variable  $z$  of complex type:6516 `z == creal(z) + cimag(z)*I`6517 **RATIONALE**

6518 None.

6519 **FUTURE DIRECTIONS**

6520 None.

6521 **SEE ALSO**6522 `carg()`, `conj()`, `cproj()`, `creal()`, the Base Definitions volume of IEEE Std 1003.1-200x, `<complex.h>`6523 **CHANGE HISTORY**

6524 First released in Issue 6. Derived from the ISO/IEC 9899:1999 standard.

6525 **NAME**

6526 clearerr — clear indicators on a stream

6527 **SYNOPSIS**

6528 #include <stdio.h>

6529 void clearerr(FILE \*stream);

6530 **DESCRIPTION**

6531 cx The functionality described on this reference page is aligned with the ISO C standard. Any  
6532 conflict between the requirements described here and the ISO C standard is unintentional. This  
6533 volume of IEEE Std 1003.1-200x defers to the ISO C standard.

6534 The *clearerr()* function shall clear the end-of-file and error indicators for the stream to which  
6535 *stream* points.

6536 **RETURN VALUE**

6537 The *clearerr()* function shall not return a value.

6538 **ERRORS**

6539 No errors are defined.

6540 **EXAMPLES**

6541 None.

6542 **APPLICATION USAGE**

6543 None.

6544 **RATIONALE**

6545 None.

6546 **FUTURE DIRECTIONS**

6547 None.

6548 **SEE ALSO**

6549 The Base Definitions volume of IEEE Std 1003.1-200x, <stdio.h>

6550 **CHANGE HISTORY**

6551 First released in Issue 1. Derived from Issue 1 of the SVID.



6552 **NAME**

6553 clock — report CPU time used

6554 **SYNOPSIS**

6555 #include &lt;time.h&gt;

6556 clock\_t clock(void);

6557 **DESCRIPTION**

6558 **CX** The functionality described on this reference page is aligned with the ISO C standard. Any  
6559 conflict between the requirements described here and the ISO C standard is unintentional. This  
6560 volume of IEEE Std 1003.1-200x defers to the ISO C standard.

6561 The *clock()* function shall return the implementation's best approximation to the processor time  
6562 used by the process since the beginning of an implementation-defined era related only to the  
6563 process invocation.

6564 **RETURN VALUE**

6565 To determine the time in seconds, the value returned by *clock()* should be divided by the value  
6566 **XSI** of the macro `CLOCKS_PER_SEC`. `CLOCKS_PER_SEC` is defined to be one million in `<time.h>`.  
6567 If the processor time used is not available or its value cannot be represented, the function shall  
6568 return the value `(clock_t)-1`.

6569 **ERRORS**

6570 No errors are defined.

6571 **EXAMPLES**

6572 None.

6573 **APPLICATION USAGE**

6574 In order to measure the time spent in a program, *clock()* should be called at the start of the  
6575 program and its return value subtracted from the value returned by subsequent calls. The value  
6576 returned by *clock()* is defined for compatibility across systems that have clocks with different  
6577 resolutions. The resolution on any particular system need not be to microsecond accuracy.

6578 The value returned by *clock()* may wrap around on some implementations. For example, on a  
6579 machine with 32-bit values for `clock_t`, it wraps after 2 147 seconds or 36 minutes.

6580 **RATIONALE**

6581 None.

6582 **FUTURE DIRECTIONS**

6583 None.

6584 **SEE ALSO**

6585 *asctime()*, *ctime()*, *difftime()*, *gmtime()*, *localtime()*, *mktime()*, *strftime()*, *strptime()*, *time()*, *utime()*,  
6586 the Base Definitions volume of IEEE Std 1003.1-200x, `<time.h>`

6587 **CHANGE HISTORY**

6588 First released in Issue 1. Derived from Issue 1 of the SVID.

6589 **NAME**6590 clock\_getcpuclockid — access a process CPU-time clock (**ADVANCED REALTIME**)6591 **SYNOPSIS**

6592 CPT #include &lt;time.h&gt;

6593 int clock\_getcpuclockid(pid\_t pid, clockid\_t \*clock\_id);

6594

6595 **DESCRIPTION**6596 The *clock\_getcpuclockid()* function shall return the clock ID of the CPU-time clock of the process  
6597 specified by *pid*. If the process described by *pid* exists and the calling process has permission,  
6598 the clock ID of this clock shall be returned in *clock\_id*.6599 If *pid* is zero, the *clock\_getcpuclockid()* function shall return the clock ID of the CPU-time clock of  
6600 the process making the call, in *clock\_id*.6601 The conditions under which one process has permission to obtain the CPU-time clock ID of  
6602 other processes are implementation-defined.6603 **RETURN VALUE**6604 Upon successful completion, *clock\_getcpuclockid()* shall return zero; otherwise, an error number  
6605 shall be returned to indicate the error.6606 **ERRORS**6607 The *clock\_getcpuclockid()* function shall fail if:6608 [EPERM] The requesting process does not have permission to access the CPU-time  
6609 clock for the process.6610 The *clock\_getcpuclockid()* function may fail if:6611 [ESRCH] No process can be found corresponding to the process specified by *pid*.6612 **EXAMPLES**

6613 None.

6614 **APPLICATION USAGE**6615 The *clock\_getcpuclockid()* function is part of the Process CPU-Time Clocks option and need not  
6616 be provided on all implementations.6617 **RATIONALE**

6618 None.

6619 **FUTURE DIRECTIONS**

6620 None.

6621 **SEE ALSO**6622 *clock\_getres()*, *timer\_create()*, the Base Definitions volume of IEEE Std 1003.1-200x, <time.h>6623 **CHANGE HISTORY**

6624 First released in Issue 6. Derived from IEEE Std 1003.1d-1999.

6625 In the SYNOPSIS, the inclusion of &lt;sys/types.h&gt; is no longer required.

## 6626 NAME

6627 clock\_getres, clock\_gettime, clock\_settime — clock and timer functions (REALTIME)

## 6628 SYNOPSIS

6629 TMR #include &lt;time.h&gt;

```
6630 int clock_getres(clockid_t clock_id, struct timespec *res);
6631 int clock_gettime(clockid_t clock_id, struct timespec *tp);
6632 int clock_settime(clockid_t clock_id, const struct timespec *tp);
6633
```

## 6634 DESCRIPTION

6635 The *clock\_getres()* function shall return the resolution of any clock. Clock resolutions are  
 6636 implementation-defined and cannot be set by a process. If the argument *res* is not NULL, the  
 6637 resolution of the specified clock shall be stored in the location pointed to by *res*. If *res* is NULL,  
 6638 the clock resolution is not returned. If the *time* argument of *clock\_settime()* is not a multiple of *res*,  
 6639 then the value is truncated to a multiple of *res*.

6640 The *clock\_gettime()* function shall return the current value *tp* for the specified clock, *clock\_id*.

6641 The *clock\_settime()* function shall set the specified clock, *clock\_id*, to the value specified by *tp*.  
 6642 Time values that are between two consecutive non-negative integer multiples of the resolution  
 6643 of the specified clock shall be truncated down to the smaller multiple of the resolution.

6644 A clock may be system-wide (that is, visible to all processes) or per-process (measuring time that  
 6645 is meaningful only within a process). All implementations shall support a *clock\_id* of  
 6646 CLOCK\_REALTIME as defined in <time.h>. This clock represents the realtime clock for the  
 6647 system. For this clock, the values returned by *clock\_gettime()* and specified by *clock\_settime()*  
 6648 represent the amount of time (in seconds and nanoseconds) since the Epoch. An implementation  
 6649 may also support additional clocks. The interpretation of time values for these clocks is  
 6650 unspecified.

6651 If the value of the CLOCK\_REALTIME clock is set via *clock\_settime()*, the new value of the clock  
 6652 shall be used to determine the time of expiration for absolute time services based upon the  
 6653 CLOCK\_REALTIME clock. This applies to the time at which armed absolute timers expire. If the  
 6654 absolute time requested at the invocation of such a time service is before the new value of the  
 6655 clock, the time service shall expire immediately as if the clock had reached the requested time  
 6656 normally.

6657 Setting the value of the CLOCK\_REALTIME clock via *clock\_settime()* shall have no effect on  
 6658 threads that are blocked waiting for a relative time service based upon this clock, including the  
 6659 *nanosleep()* function; nor on the expiration of relative timers based upon this clock.  
 6660 Consequently, these time services shall expire when the requested relative interval elapses,  
 6661 independently of the new or old value of the clock.

6662 MON If the Monotonic Clock option is supported, all implementations shall support a *clock\_id* of  
 6663 CLOCK\_MONOTONIC defined in <time.h>. This clock represents the monotonic clock for the  
 6664 system. For this clock, the value returned by *clock\_gettime()* represents the amount of time (in  
 6665 seconds and nanoseconds) since an unspecified point in the past (for example, system start-up  
 6666 time, or the Epoch). This point does not change after system start-up time. The value of the  
 6667 CLOCK\_MONOTONIC clock cannot be set via *clock\_settime()*. This function shall fail if it is  
 6668 invoked with a *clock\_id* argument of CLOCK\_MONOTONIC.

6669 The effect of setting a clock via *clock\_settime()* on armed per-process timers associated with a  
 6670 clock other than CLOCK\_REALTIME is implementation-defined.

6671 CS If the value of the CLOCK\_REALTIME clock is set via *clock\_settime()*, the new value of the clock  
 6672 shall be used to determine the time at which the system shall awaken a thread blocked on an

6673 absolute *clock\_nanosleep()* call based upon the CLOCK\_REALTIME clock. If the absolute time  
 6674 requested at the invocation of such a time service is before the new value of the clock, the call  
 6675 shall return immediately as if the clock had reached the requested time normally.

6676 Setting the value of the CLOCK\_REALTIME clock via *clock\_settime()* shall have no effect on any  
 6677 thread that is blocked on a relative *clock\_nanosleep()* call. Consequently, the call shall return  
 6678 when the requested relative interval elapses, independently of the new or old value of the clock.

6679 The appropriate privilege to set a particular clock is implementation-defined.

6680 CPT If `_POSIX_CPUTIME` is defined, implementations shall support clock ID values obtained by  
 6681 invoking *clock\_getcpuclockid()*, which represent the CPU-time clock of a given process.  
 6682 Implementations shall also support the special `clockid_t` value  
 6683 `CLOCK_PROCESS_CPUTIME_ID`, which represents the CPU-time clock of the calling process  
 6684 when invoking one of the *clock\_\**() or *timer\_\**() functions. For these clock IDs, the values  
 6685 returned by *clock\_gettime()* and specified by *clock\_settime()* represent the amount of execution  
 6686 time of the process associated with the clock. Changing the value of a CPU-time clock via  
 6687 *clock\_settime()* shall have no effect on the behavior of the sporadic server scheduling policy (see  
 6688 **Scheduling Policies** (on page 494)).

6689 TCT If `_POSIX_THREAD_CPUTIME` is defined, implementations shall support clock ID values  
 6690 obtained by invoking *pthread\_getcpuclockid()*, which represent the CPU-time clock of a given  
 6691 thread. Implementations shall also support the special `clockid_t` value  
 6692 `CLOCK_THREAD_CPUTIME_ID`, which represents the CPU-time clock of the calling thread  
 6693 when invoking one of the *clock\_\**() or *timer\_\**() functions. For these clock IDs, the values  
 6694 returned by *clock\_gettime()* and specified by *clock\_settime()* shall represent the amount of  
 6695 execution time of the thread associated with the clock. Changing the value of a CPU-time clock  
 6696 via *clock\_settime()* shall have no effect on the behavior of the sporadic server scheduling policy  
 6697 (see **Scheduling Policies** (on page 494)).

#### 6698 RETURN VALUE

6699 A return value of 0 shall indicate that the call succeeded. A return value of -1 shall indicate that  
 6700 an error occurred, and *errno* shall be set to indicate the error.

#### 6701 ERRORS

6702 The *clock\_getres()*, *clock\_gettime()*, and *clock\_settime()* functions shall fail if:

6703 [EINVAL] The *clock\_id* argument does not specify a known clock.

6704 The *clock\_settime()* function shall fail if:

6705 [EINVAL] The *tp* argument to *clock\_settime()* is outside the range for the given clock ID.

6706 [EINVAL] The *tp* argument specified a nanosecond value less than zero or greater than  
 6707 or equal to 1 000 million.

6708 MON [EINVAL] The value of the *clock\_id* argument is `CLOCK_MONOTONIC`.

6709 The *clock\_settime()* function may fail if:

6710 [EPERM] The requesting process does not have the appropriate privilege to set the  
 6711 specified clock.

6712 **EXAMPLES**

6713 None.

6714 **APPLICATION USAGE**

6715 These functions are part of the Timers option and need not be available on all implementations.

6716 Note that the absolute value of the monotonic clock is meaningless (because its origin is  
 6717 arbitrary), and thus there is no need to set it. Furthermore, realtime applications can rely on the  
 6718 fact that the value of this clock is never set and, therefore, that time intervals measured with this  
 6719 clock will not be affected by calls to *clock\_settime()*.

6720 **RATIONALE**

6721 None.

6722 **FUTURE DIRECTIONS**

6723 None.

6724 **SEE ALSO**

6725 *clock\_getcpuclockid()*, *clock\_nanosleep()*, *ctime()*, *mq\_timedreceive()*, *mq\_timedsend()*, *nanosleep()*,  
 6726 *pthread\_mutex\_timedlock()*, *sem\_timedwait()*, *time()*, *timer\_create()*, *timer\_getoverrun()*, the Base  
 6727 Definitions volume of IEEE Std 1003.1-200x, <**time.h**>

6728 **CHANGE HISTORY**

6729 First released in Issue 5. Included for alignment with the POSIX Realtime Extension.

6730 **Issue 6**

6731 The [ENOSYS] error condition has been removed as stubs need not be provided if an  
 6732 implementation does not support the Timers option.

6733 The APPLICATION USAGE section is added.

6734 The following changes were made to align with the IEEE P1003.1a draft standard:

- 6735 • Clarification is added of the effect of resetting the clock resolution.

6736 CPU-time clocks and the *clock\_getcpuclockid()* function are added for alignment with  
 6737 IEEE Std 1003.1d-1999.

6738 The following changes are added for alignment with IEEE Std 1003.1j-2000:

- 6739 • The DESCRIPTION is updated as follows:

- 6740 — The value returned by *clock\_gettime()* for CLOCK\_MONOTONIC is specified.

- 6741 — *clock\_settime()* failing for CLOCK\_MONOTONIC is specified.

- 6742 — The effects of *clock\_settime()* on the *clock\_nanosleep()* function with respect to  
 6743 CLOCK\_REALTIME is specified.

- 6744 • An [EINVAL] error is added to the ERRORS section, indicating that *clock\_settime()* fails for  
 6745 CLOCK\_MONOTONIC.

- 6746 • The APPLICATION USAGE section notes that the CLOCK\_MONOTONIC clock need not  
 6747 and shall not be set by *clock\_settime()* since the absolute value of the CLOCK\_MONOTONIC  
 6748 clock is meaningless.

- 6749 • The *clock\_nanosleep()*, *mq\_timedreceive()*, *mq\_timedsend()*, *pthread\_mutex\_timedlock()*,  
 6750 *sem\_timedwait()*, *timer\_create()*, and *timer\_settime()* functions are added to the SEE ALSO  
 6751 section.

## 6752 NAME

6753 clock\_nanosleep — high resolution sleep with specifiable clock (**ADVANCED REALTIME**)

## 6754 SYNOPSIS

```
6755 cs #include <time.h>
```

```
6756 int clock_nanosleep(clockid_t clock_id, int flags,  
6757 const struct timespec *rqtp, struct timespec *rmtp);  
6758
```

## 6759 DESCRIPTION

6760 If the flag `TIMER_ABSTIME` is not set in the *flags* argument, the *clock\_nanosleep()* function shall  
6761 cause the current thread to be suspended from execution until either the time interval specified  
6762 by the *rqtp* argument has elapsed, or a signal is delivered to the calling thread and its action is to  
6763 invoke a signal-catching function, or the process is terminated. The clock used to measure the  
6764 time shall be the clock specified by *clock\_id*.

6765 If the flag `TIMER_ABSTIME` is set in the *flags* argument, the *clock\_nanosleep()* function shall  
6766 cause the current thread to be suspended from execution until either the time value of the clock  
6767 specified by *clock\_id* reaches the absolute time specified by the *rqtp* argument, or a signal is  
6768 delivered to the calling thread and its action is to invoke a signal-catching function, or the  
6769 process is terminated. If, at the time of the call, the time value specified by *rqtp* is less than or  
6770 equal to the time value of the specified clock, then *clock\_nanosleep()* shall return immediately  
6771 and the calling process shall not be suspended.

6772 The suspension time caused by this function may be longer than requested because the  
6773 argument value is rounded up to an integer multiple of the sleep resolution, or because of the  
6774 scheduling of other activity by the system. But, except for the case of being interrupted by a  
6775 signal, the suspension time for the relative *clock\_nanosleep()* function (that is, with the  
6776 `TIMER_ABSTIME` flag not set) shall not be less than the time interval specified by *rqtp*, as  
6777 measured by the corresponding clock. The suspension for the absolute *clock\_nanosleep()* function  
6778 (that is, with the `TIMER_ABSTIME` flag set) shall be in effect at least until the value of the  
6779 corresponding clock reaches the absolute time specified by *rqtp*, except for the case of being  
6780 interrupted by a signal.

6781 The use of the *clock\_nanosleep()* function shall have no effect on the action or blockage of any  
6782 signal.

6783 The *clock\_nanosleep()* function shall fail if the *clock\_id* argument refers to the CPU-time clock of  
6784 the calling thread. It is unspecified if *clock\_id* values of other CPU-time clocks are allowed.

## 6785 RETURN VALUE

6786 If the *clock\_nanosleep()* function returns because the requested time has elapsed, its return value  
6787 shall be zero.

6788 If the *clock\_nanosleep()* function returns because it has been interrupted by a signal, it shall return  
6789 the corresponding error value. For the relative *clock\_nanosleep()* function, if the *rmtp* argument is  
6790 non-NULL, the **timespec** structure referenced by it shall be updated to contain the amount of  
6791 time remaining in the interval (the requested time minus the time actually slept). If the *rmtp*  
6792 argument is NULL, the remaining time is not returned. The absolute *clock\_nanosleep()* function  
6793 has no effect on the structure referenced by *rmtp*.

6794 If *clock\_nanosleep()* fails, it shall return the corresponding error value.

6795 **ERRORS**6796 The *clock\_nanosleep()* function shall fail if:

- 6797 [EINTR] The *clock\_nanosleep()* function was interrupted by a signal.
- 6798 [EINVAL] The *rntp* argument specified a nanosecond value less than zero or greater than  
6799 or equal to 1 000 million; or the `TIMER_ABSTIME` flag was specified in *flags*  
6800 and the *rntp* argument is outside the range for the clock specified by *clock\_id*;  
6801 or the *clock\_id* argument does not specify a known clock, or specifies the  
6802 CPU-time clock of the calling thread.
- 6803 [ENOTSUP] The *clock\_id* argument specifies a clock for which *clock\_nanosleep()* is not  
6804 supported, such as a CPU-time clock.

6805 **EXAMPLES**

6806 None.

6807 **APPLICATION USAGE**

6808 Calling *clock\_nanosleep()* with the value `TIMER_ABSTIME` not set in the *flags* argument and with  
6809 a *clock\_id* of `CLOCK_REALTIME` is equivalent to calling *nanosleep()* with the same *rntp* and *rntp*  
6810 arguments.

6811 **RATIONALE**

6812 The *nanosleep()* function specifies that the system-wide clock `CLOCK_REALTIME` is used to  
6813 measure the elapsed time for this time service. However, with the introduction of the monotonic  
6814 clock `CLOCK_MONOTONIC` a new relative sleep function is needed to allow an application to  
6815 take advantage of the special characteristics of this clock.

6816 There are many applications in which a process needs to be suspended and then activated  
6817 multiple times in a periodic way; for example, to poll the status of a non-interrupting device or  
6818 to refresh a display device. For these cases, it is known that precise periodic activation cannot be  
6819 achieved with a relative *sleep()* or *nanosleep()* function call. Suppose, for example, a periodic  
6820 process that is activated at time  $T_0$ , executes for a while, and then wants to suspend itself until  
6821 time  $T_0+T$ , the period being  $T$ . If this process wants to use the *nanosleep()* function, it must first  
6822 call *clock\_gettime()* to get the current time, then calculate the difference between the current time  
6823 and  $T_0+T$  and, finally, call *nanosleep()* using the computed interval. However, the process could  
6824 be preempted by a different process between the two function calls, and in this case the interval  
6825 computed would be wrong; the process would wake up later than desired. This problem would  
6826 not occur with the absolute *clock\_nanosleep()* function, since only one function call would be  
6827 necessary to suspend the process until the desired time. In other cases, however, a relative sleep  
6828 is needed, and that is why both functionalities are required.

6829 Although it is possible to implement periodic processes using the timers interface, this  
6830 implementation would require the use of signals, and the reservation of some signal numbers. In  
6831 this regard, the reasons for including an absolute version of the *clock\_nanosleep()* function in  
6832 IEEE Std 1003.1-200x are the same as for the inclusion of the relative *nanosleep()*.

6833 It is also possible to implement precise periodic processes using *pthread\_cond\_timedwait()*, in  
6834 which an absolute timeout is specified that takes effect if the condition variable involved is  
6835 never signaled. However, the use of this interface is unnatural, and involves performing other  
6836 operations on mutexes and condition variables that imply an unnecessary overhead.  
6837 Furthermore, *pthread\_cond\_timedwait()* is not available in implementations that do not support  
6838 threads.

6839 Although the interface of the relative and absolute versions of the new high resolution sleep  
6840 service is the same *clock\_nanosleep()* function, the *rntp* argument is only used in the relative  
6841 sleep. This argument is needed in the relative *clock\_nanosleep()* function to reissue the function

6842 call if it is interrupted by a signal, but it is not needed in the absolute *clock\_nanosleep()* function  
6843 call; if the call is interrupted by a signal, the absolute *clock\_nanosleep()* function can be invoked  
6844 again with the same *rqt* argument used in the interrupted call.

6845 **FUTURE DIRECTIONS**

6846 None.

6847 **SEE ALSO**

6848 *clock\_getres()*, *nanosleep()*, *pthread\_cond\_timedwait()*, *sleep()*, the Base Definitions volume of  
6849 IEEE Std 1003.1-200x, <**time.h**>

6850 **CHANGE HISTORY**

6851 First released in Issue 6. Derived from IEEE Std 1003.1j-2000.



6852 **NAME**6853 clock\_settime — clock and timer functions (**REALTIME**)6854 **SYNOPSIS**

6855 TMR #include &lt;time.h&gt;

6856 int clock\_settime(clockid\_t *clock\_id*, const struct timespec \**tp*);

6857

6858 **DESCRIPTION**6859 Refer to *clock\_getres()*.

6860 **NAME**6861 `clog, clogf, clogl` — complex natural logarithm functions6862 **SYNOPSIS**6863 `#include <complex.h>`6864 `double complex clog(double complex z);`6865 `float complex clogf(float complex z);`6866 `long double complex clogl(long double complex z);`6867 **DESCRIPTION**

6868 **CX** The functionality described on this reference page is aligned with the ISO C standard. Any  
6869 conflict between the requirements described here and the ISO C standard is unintentional. This  
6870 volume of IEEE Std 1003.1-200x defers to the ISO C standard.

6871 These functions shall compute the complex natural (base  $e$ ) logarithm of  $z$ , with a branch cut  
6872 along the negative real axis.

6873 **RETURN VALUE**

6874 These functions shall return the complex natural logarithm value, in the range of a strip  
6875 mathematically unbounded along the real axis and in the interval  $[-i\pi, +i\pi]$  along the imaginary  
6876 axis.

6877 **ERRORS**

6878 No errors are defined.

6879 **EXAMPLES**

6880 None.

6881 **APPLICATION USAGE**

6882 None.

6883 **RATIONALE**

6884 None.

6885 **FUTURE DIRECTIONS**

6886 None.

6887 **SEE ALSO**6888 `cexp()`, the Base Definitions volume of IEEE Std 1003.1-200x, `<complex.h>`6889 **CHANGE HISTORY**

6890 First released in Issue 6. Derived from the ISO/IEC 9899:1999 standard.

6891 **NAME**

6892           close — close a file descriptor

6893 **SYNOPSIS**

6894           #include &lt;unistd.h&gt;

6895           int close(int *fdes*);6896 **DESCRIPTION**

6897       The *close()* function shall deallocate the file descriptor indicated by *fdes*. To deallocate means  
 6898       to make the file descriptor available for return by subsequent calls to *open()* or other functions  
 6899       that allocate file descriptors. All outstanding record locks owned by the process on the file  
 6900       associated with the file descriptor shall be removed (that is, unlocked).

6901       If *close()* is interrupted by a signal that is to be caught, it shall return  $-1$  with *errno* set to [EINTR]  
 6902       and the state of *fdes* is unspecified. If an I/O error occurred while reading from or writing to the  
 6903       file system during *close()*, it may return  $-1$  with *errno* set to [EIO]; if this error is returned, the  
 6904       state of *fdes* is unspecified.

6905       When all file descriptors associated with a pipe or FIFO special file are closed, any data  
 6906       remaining in the pipe or FIFO shall be discarded.

6907       When all file descriptors associated with an open file description have been closed the open file  
 6908       description shall be freed.

6909       If the link count of the file is 0, when all file descriptors associated with the file are closed, the  
 6910       space occupied by the file shall be freed and the file shall no longer be accessible.

6911 **XSR**       If a STREAMS-based *fdes* is closed and the calling process was previously registered to receive  
 6912       a SIGPOLL signal for events associated with that STREAM, the calling process shall be  
 6913       unregistered for events associated with the STREAM. The last *close()* for a STREAM shall cause  
 6914       the STREAM associated with *fdes* to be dismantled. If O\_NONBLOCK is not set and there have  
 6915       been no signals posted for the STREAM, and if there is data on the module's write queue, *close()*  
 6916       shall wait for an unspecified time (for each module and driver) for any output to drain before  
 6917       dismantling the STREAM. The time delay can be changed via an I\_SETCLTIME *ioctl()* request. If  
 6918       the O\_NONBLOCK flag is set, or if there are any pending signals, *close()* shall not wait for  
 6919       output to drain, and shall dismantle the STREAM immediately.

6920       If the implementation supports STREAMS-based pipes, and *fdes* is associated with one end of a  
 6921       pipe, the last *close()* shall cause a hangup to occur on the other end of the pipe. In addition, if the  
 6922       other end of the pipe has been named by *fattach()*, then the last *close()* shall force the named end  
 6923       to be detached by *fdetach()*. If the named end has no open file descriptors associated with it and  
 6924       gets detached, the STREAM associated with that end shall also be dismantled.

6925 **XSI**       If *fdes* refers to the master side of a pseudo-terminal, and this is the last close, a SIGHUP signal  
 6926       shall be sent to the process group, if any, for which the slave side of the pseudo-terminal is the  
 6927       controlling terminal. It is unspecified whether closing the master side of the pseudo-terminal  
 6928       flushes all queued input and output.

6929 **XSR**       If *fdes* refers to the slave side of a STREAMS-based pseudo-terminal, a zero-length message  
 6930       may be sent to the master.

6931 **AIO**       When there is an outstanding cancelable asynchronous I/O operation against *fdes* when *close()*  
 6932       is called, that I/O operation may be canceled. An I/O operation that is not canceled completes  
 6933       as if the *close()* operation had not yet occurred. All operations that are not canceled shall  
 6934       complete as if the *close()* blocked until the operations completed. The *close()* operation itself  
 6935       need not block awaiting such I/O completion. Whether any I/O operation is canceled, and  
 6936       which I/O operation may be canceled upon *close()*, is implementation-defined.

6937 MF|SHM If a shared memory object or a memory mapped file remains referenced at the last close (that is,  
6938 a process has it mapped), then the entire contents of the memory object shall persist until the  
6939 memory object becomes unreferenced. If this is the last close of a shared memory object or a  
6940 memory mapped file and the close results in the memory object becoming unreferenced, and the  
6941 memory object has been unlinked, then the memory object shall be removed.

6942 If *fdes* refers to a socket, *close()* shall cause the socket to be destroyed. If the socket is in  
6943 connection-mode, and the `SO_LINGER` option is set for the socket with non-zero linger time,  
6944 and the socket has untransmitted data, then *close()* shall block for up to the current linger  
6945 interval until all data is transmitted.

#### 6946 RETURN VALUE

6947 Upon successful completion, 0 shall be returned; otherwise, -1 shall be returned and *errno* set to  
6948 indicate the error.

#### 6949 ERRORS

6950 The *close()* function shall fail if:

6951 [EBADF] The *fdes* argument is not a valid file descriptor.

6952 [EINTR] The *close()* function was interrupted by a signal.

6953 The *close()* function may fail if:

6954 [EIO] An I/O error occurred while reading from or writing to the file system.

#### 6955 EXAMPLES

##### 6956 Reassigning a File Descriptor

6957 The following example closes the file descriptor associated with standard output for the current  
6958 process, re-assigns standard output to a new file descriptor, and closes the original file  
6959 descriptor to clean up. This example assumes that the file descriptor 0 (which is the descriptor  
6960 for standard input) is not closed.

```
6961 #include <unistd.h>  
6962 ...  
6963 int pfd;  
6964 ...  
6965 close(1);  
6966 dup(pfd);  
6967 close(pfd);  
6968 ...
```

6969 Incidentally, this is exactly what could be achieved using:

```
6970 dup2(pfd, 1);  
6971 close(pfd);
```

##### 6972 Closing a File Descriptor

6973 In the following example, *close()* is used to close a file descriptor after an unsuccessful attempt is  
6974 made to associate that file descriptor with a stream.

```
6975 #include <stdio.h>  
6976 #include <unistd.h>  
6977 #include <stdlib.h>
```

```

6978     #define LOCKFILE "/etc/ptmp"
6979     ...
6980     int pfd;
6981     FILE *fpfd;
6982     ...
6983     if ((fpfd = fdopen (pfd, "w")) == NULL) {
6984         close(pfd);
6985         unlink(LOCKFILE);
6986         exit(1);
6987     }
6988     ...

```

#### 6989 APPLICATION USAGE

6990 An application that had used the *stdio* routine *fopen()* to open a file should use the  
 6991 corresponding *fclose()* routine rather than *close()*. Once a file is closed, the file descriptor no  
 6992 longer exists, since the integer corresponding to it no longer refers to a file.

#### 6993 RATIONALE

6994 The use of interruptible device close routines should be discouraged to avoid problems with the  
 6995 implicit closes of file descriptors by *exec* and *exit()*. This volume of IEEE Std 1003.1-200x only  
 6996 intends to permit such behavior by specifying the [EINTR] error condition.

#### 6997 FUTURE DIRECTIONS

6998 None.

#### 6999 SEE ALSO

7000 *fattach()*, *fclose()*, *fdetach()*, *fopen()*, *ioctl()*, *open()*, the Base Definitions volume of  
 7001 IEEE Std 1003.1-200x, <*unistd.h*>, Section 2.6 (on page 488)

#### 7002 CHANGE HISTORY

7003 First released in Issue 1. Derived from Issue 1 of the SVID.

#### 7004 Issue 5

7005 The DESCRIPTION is updated for alignment with the POSIX Realtime Extension.

#### 7006 Issue 6

7007 The DESCRIPTION related to a STREAMS-based file or pseudo-terminal is marked as part of the  
 7008 XSI STREAMS Option Group.

7009 The following new requirements on POSIX implementations derive from alignment with the  
 7010 Single UNIX Specification:

- 7011 • The [EIO] error condition is added as an optional error.
- 7012 • The DESCRIPTION is updated to describe the state of the *fildev* file descriptor as unspecified  
 7013 if an I/O error occurs and an [EIO] error condition is returned.

7014 Text referring to sockets is added to the DESCRIPTION.

7015 The DESCRIPTION is updated for alignment with IEEE Std 1003.1j-2000 by specifying that  
 7016 shared memory objects and memory mapped files (and not typed memory objects) are the types  
 7017 of memory objects to which the paragraph on last closes applies.

7018 **NAME**

7019        closedir — close a directory stream

7020 **SYNOPSIS**

7021        #include &lt;dirent.h&gt;

7022        int closedir(DIR \*dirp);

7023 **DESCRIPTION**

7024        The *closedir()* function shall close the directory stream referred to by the argument *dirp*. Upon  
7025        return, the value of *dirp* may no longer point to an accessible object of the type **DIR**. If a file  
7026        descriptor is used to implement type **DIR**, that file descriptor shall be closed.

7027 **RETURN VALUE**

7028        Upon successful completion, *closedir()* shall return 0; otherwise, -1 shall be returned and *errno*  
7029        set to indicate the error.

7030 **ERRORS**7031        The *closedir()* function may fail if:7032        [EBADF]        The *dirp* argument does not refer to an open directory stream.7033        [EINTR]        The *closedir()* function was interrupted by a signal.7034 **EXAMPLES**7035        **Closing a Directory Stream**7036        The following program fragment demonstrates how the *closedir()* function is used.

```
7037        ...  
7038        DIR *dir;  
7039        struct dirent *dp;  
7040        ...  
7041        if ((dir = opendir (".")) == NULL) {  
7042        ...  
7043        }  
7044        while ((dp = readdir (dir)) != NULL) {  
7045        ...  
7046        }  
7047        closedir(dir);  
7048        ...
```

7049 **APPLICATION USAGE**

7050        None.

7051 **RATIONALE**

7052        None.

7053 **FUTURE DIRECTIONS**

7054        None.

7055 **SEE ALSO**7056        *opendir()*, the Base Definitions volume of IEEE Std 1003.1-200x, <dirent.h>

7057 **CHANGE HISTORY**

7058 First released in Issue 2.

7059 **Issue 6**7060 In the SYNOPSIS, the optional include of the `<sys/types.h>` header is removed.7061 The following new requirements on POSIX implementations derive from alignment with the  
7062 Single UNIX Specification:

- 7063 • The requirement to include `<sys/types.h>` has been removed. Although `<sys/types.h>` was  
7064 required for conforming implementations of previous POSIX specifications, it was not  
7065 required for UNIX applications.
- 7066 • The [EINTR] error condition is added as an optional error condition.

7067 **NAME**

7068 closelog, openlog, setlogmask, syslog — control system log

7069 **SYNOPSIS**

```
7070 xSI #include <syslog.h>
7071
7072 void closelog(void);
7073 void openlog(const char *ident, int logopt, int facility);
7074 int setlogmask(int maskpri);
7075 void syslog(int priority, const char *message, ... /* arguments */);
```

7076 **DESCRIPTION**

7077 The *syslog()* function shall send a message to an implementation-defined logging facility, which  
 7078 may log it in an implementation-defined system log, write it to the system console, forward it to  
 7079 a list of users, or forward it to the logging facility on another host over the network. The logged  
 7080 message shall include a message header and a message body. The message header contains at  
 7081 least a timestamp and a tag string.

7082 The message body is generated from the *message* and following arguments in the same manner  
 7083 as if these were arguments to *printf()*, except that the additional conversion specification *%m*  
 7084 shall be recognized; it shall convert no arguments, shall cause the output of the error message  
 7085 string associated with the value of *errno* on entry to *syslog()*, and may be mixed with argument  
 7086 specifications of the "*%n\$*" form. If a complete conversion specification with the *m* conversion  
 7087 specifier character is not just *%m*, the behavior is undefined. A trailing *<newline>* may be added  
 7088 if needed.

7089 Values of the *priority* argument are formed by OR'ing together a severity level value and an  
 7090 optional facility value. If no facility value is specified, the current default facility value is used.

7091 Possible values of severity level include:

7092	LOG_EMERG	A panic condition.
7093	LOG_ALERT	A condition that should be corrected immediately, such as a corrupted system database.
7094		
7095	LOG_CRIT	Critical conditions, such as hard device errors.
7096	LOG_ERR	Errors.
7097	LOG_WARNING	
7098		Warning messages.
7099	LOG_NOTICE	Conditions that are not error conditions, but that may require special handling.
7100		
7101	LOG_INFO	Informational messages.
7102	LOG_DEBUG	Messages that contain information normally of use only when debugging a program.
7103		

7104 The facility indicates the application or system component generating the message. Possible  
 7105 facility values include:

7106	LOG_USER	Messages generated by arbitrary processes. This is the default facility identifier if none is specified.
7107		
7108	LOG_LOCAL0	Reserved for local use.



7109	LOG_LOCAL1	Reserved for local use.
7110	LOG_LOCAL2	Reserved for local use.
7111	LOG_LOCAL3	Reserved for local use.
7112	LOG_LOCAL4	Reserved for local use.
7113	LOG_LOCAL5	Reserved for local use.
7114	LOG_LOCAL6	Reserved for local use.
7115	LOG_LOCAL7	Reserved for local use.
7116	The <i>openlog()</i> function shall set process attributes that affect subsequent calls to <i>syslog()</i> . The	
7117	<i>ident</i> argument is a string that is prepended to every message. The <i>logopt</i> argument indicates	
7118	logging options. Values for <i>logopt</i> are constructed by a bitwise-inclusive OR of zero or more of	
7119	the following:	
7120	LOG_PID	Log the process ID with each message. This is useful for identifying specific
7121		processes.
7122	LOG_CONS	Write messages to the system console if they cannot be sent to the logging
7123		facility. The <i>syslog()</i> function ensures that the process does not acquire the
7124		console as a controlling terminal in the process of writing the message.
7125	LOG_NDELAY	Open the connection to the logging facility immediately. Normally the open is
7126		delayed until the first message is logged. This is useful for programs that need
7127		to manage the order in which file descriptors are allocated.
7128	LOG_ODELAY	Delay open until <i>syslog()</i> is called.
7129	LOG_NOWAIT	Do not wait for child processes that may have been created during the course
7130		of logging the message. This option should be used by processes that enable
7131		notification of child termination using SIGCHLD, since <i>syslog()</i> may
7132		otherwise block waiting for a child whose exit status has already been
7133		collected.
7134	The <i>facility</i> argument encodes a default facility to be assigned to all messages that do not have	
7135	an explicit facility already encoded. The initial default facility is LOG_USER.	
7136	The <i>openlog()</i> and <i>syslog()</i> functions may allocate a file descriptor. It is not necessary to call	
7137	<i>openlog()</i> prior to calling <i>syslog()</i> .	
7138	The <i>closelog()</i> function shall close any open file descriptors allocated by previous calls to	
7139	<i>openlog()</i> or <i>syslog()</i> .	
7140	The <i>setlogmask()</i> function shall set the log priority mask for the current process to <i>maskpri</i> and	
7141	return the previous mask. If the <i>maskpri</i> argument is 0, the current log mask is not modified.	
7142	Calls by the current process to <i>syslog()</i> with a priority not set in <i>maskpri</i> shall be rejected. The	
7143	default log mask allows all priorities to be logged. A call to <i>openlog()</i> is not required prior to	
7144	calling <i>setlogmask()</i> .	
7145	Symbolic constants for use as values of the <i>logopt</i> , <i>facility</i> , <i>priority</i> , and <i>maskpri</i> arguments are	
7146	defined in the <syslog.h> header.	
7147	<b>RETURN VALUE</b>	
7148	The <i>setlogmask()</i> function shall return the previous log priority mask. The <i>closelog()</i> , <i>openlog()</i> ,	
7149	and <i>syslog()</i> functions shall not return a value.	

7150 **ERRORS**

7151       None.

7152 **EXAMPLES**7153       **Using openlog()**

7154       The following example causes subsequent calls to *syslog()* to log the process ID with each message, and to write messages to the system console if they cannot be sent to the logging facility.

```
7157       #include <syslog.h>
7158       char *ident = "Process demo";
7159       int logopt = LOG_PID | LOG_CONS;
7160       int facility = LOG_USER;
7161       ...
7162       openlog(ident, logopt, facility);
```

7163       **Using setlogmask()**

7164       The following example causes subsequent calls to *syslog()* to accept error messages or messages generated by arbitrary processes, and to reject all other messages.

```
7166       #include <syslog.h>
7167       int result;
7168       int mask = LOG_MASK (LOG_ERR | LOG_USER);
7169       ...
7170       result = setlogmask(mask);
```

7171       **Using syslog**

7172       The following example sends the message "This is a message" to the default logging facility, marking the message as an error message generated by random processes.

```
7174       #include <syslog.h>
7175       char *message = "This is a message";
7176       int priority = LOG_ERR | LOG_USER;
7177       ...
7178       syslog(priority, message);
```

7179 **APPLICATION USAGE**

7180       None.

7181 **RATIONALE**

7182       None.

7183 **FUTURE DIRECTIONS**

7184       None.

7185 **SEE ALSO**7186       *printf()*, the Base Definitions volume of IEEE Std 1003.1-200x, <syslog.h>

7187 **CHANGE HISTORY**

7188           First released in Issue 4, Version 2.

7189 **Issue 5**

7190           Moved from X/OPEN UNIX extension to BASE.

7191 **NAME**

7192 confstr — get configurable variables

7193 **SYNOPSIS**

7194 #include &lt;unistd.h&gt;

7195 size\_t confstr(int name, char \*buf, size\_t len);

7196 **DESCRIPTION**7197 The *confstr()* function shall return configuration-defined string values. Its use and purpose are |  
7198 similar to *sysconf()*, but it is used where string values rather than numeric values are returned. |7199 The *name* argument represents the system variable to be queried. The implementation shall  
7200 support the following name values, defined in <unistd.h>. It may support others:

7201 \_CS\_PATH  
 7202 \_CS\_POSIX\_V6\_ILP32\_OFF32\_CFLAGS  
 7203 \_CS\_POSIX\_V6\_ILP32\_OFF32\_LDFLAGS  
 7204 \_CS\_POSIX\_V6\_ILP32\_OFF32\_LIBS  
 7205 \_CS\_POSIX\_V6\_ILP32\_OFF32\_LINTFLAGS  
 7206 \_CS\_POSIX\_V6\_ILP32\_OFFBIG\_CFLAGS  
 7207 \_CS\_POSIX\_V6\_ILP32\_OFFBIG\_LDFLAGS  
 7208 \_CS\_POSIX\_V6\_ILP32\_OFFBIG\_LIBS  
 7209 \_CS\_POSIX\_V6\_ILP32\_OFFBIG\_LINTFLAGS  
 7210 \_CS\_POSIX\_V6\_LP64\_OFF64\_CFLAGS  
 7211 \_CS\_POSIX\_V6\_LP64\_OFF64\_LDFLAGS  
 7212 \_CS\_POSIX\_V6\_LP64\_OFF64\_LIBS  
 7213 \_CS\_POSIX\_V6\_LP64\_OFF64\_LINTFLAGS  
 7214 \_CS\_POSIX\_V6\_LPBIG\_OFFBIG\_CFLAGS  
 7215 \_CS\_POSIX\_V6\_LPBIG\_OFFBIG\_LDFLAGS  
 7216 \_CS\_POSIX\_V6\_LPBIG\_OFFBIG\_LIBS  
 7217 \_CS\_POSIX\_V6\_LPBIG\_OFFBIG\_LINTFLAGS  
 7218 XSI CS\_XBS5\_ILP32\_OFF32\_CFLAGS (LEGACY)  
 7219 CS\_XBS5\_ILP32\_OFF32\_LDFLAGS (LEGACY)  
 7220 CS\_XBS5\_ILP32\_OFF32\_LIBS (LEGACY)  
 7221 CS\_XBS5\_ILP32\_OFF32\_LINTFLAGS (LEGACY)  
 7222 CS\_XBS5\_ILP32\_OFFBIG\_CFLAGS (LEGACY)  
 7223 CS\_XBS5\_ILP32\_OFFBIG\_LDFLAGS (LEGACY)  
 7224 CS\_XBS5\_ILP32\_OFFBIG\_LIBS (LEGACY)  
 7225 CS\_XBS5\_ILP32\_OFFBIG\_LINTFLAGS (LEGACY)  
 7226 CS\_XBS5\_LP64\_OFF64\_CFLAGS (LEGACY)  
 7227 CS\_XBS5\_LP64\_OFF64\_LDFLAGS (LEGACY)  
 7228 CS\_XBS5\_LP64\_OFF64\_LIBS (LEGACY)  
 7229 CS\_XBS5\_LP64\_OFF64\_LINTFLAGS (LEGACY)  
 7230 CS\_XBS5\_LPBIG\_OFFBIG\_CFLAGS (LEGACY)  
 7231 CS\_XBS5\_LPBIG\_OFFBIG\_LDFLAGS (LEGACY)  
 7232 CS\_XBS5\_LPBIG\_OFFBIG\_LIBS (LEGACY)  
 7233 CS\_XBS5\_LPBIG\_OFFBIG\_LINTFLAGS (LEGACY)

7234

7235 If *len* is not 0, and if *name* has a configuration-defined value, *confstr()* shall copy that value into  
7236 the *len*-byte buffer pointed to by *buf*. If the string to be returned is longer than *len* bytes,  
7237 including the terminating null, then *confstr()* shall truncate the string to *len*−1 bytes and null-  
7238 terminate the result. The application can detect that the string was truncated by comparing the  
7239 value returned by *confstr()* with *len*.

7240 If *len* is 0 and *buf* is a null pointer, then *confstr()* shall still return the integer value as defined  
 7241 below, but shall not return a string. If *len* is 0 but *buf* is not a null pointer, the result is  
 7242 unspecified.

7243 If the implementation supports the Shell option, the string stored in *buf* after a call to:

```
7244 confstr(_CS_PATH, buf, sizeof(buf))
```

7245 can be used as a value of the *PATH* environment variable that accesses all of the standard  
 7246 utilities of IEEE Std 1003.1-200x, if the return value is less than or equal to *sizeof(buf)*.

#### 7247 RETURN VALUE

7248 If *name* has a configuration-defined value, *confstr()* shall return the size of buffer that would be  
 7249 needed to hold the entire configuration-defined value including the terminating null. If this  
 7250 return value is greater than *len*, the string returned in *buf* is truncated.

7251 If *name* is invalid, *confstr()* shall return 0 and set *errno* to indicate the error.

7252 If *name* does not have a configuration-defined value, *confstr()* shall return 0 and leave *errno*  
 7253 unchanged.

#### 7254 ERRORS

7255 The *confstr()* function shall fail if:

7256 [EINVAL] The value of the *name* argument is invalid.

#### 7257 EXAMPLES

7258 None.

#### 7259 APPLICATION USAGE

7260 An application can distinguish between an invalid *name* parameter value and one that  
 7261 corresponds to a configurable variable that has no configuration-defined value by checking if  
 7262 *errno* is modified. This mirrors the behavior of *sysconf()*.

7263 The original need for this function was to provide a way of finding the configuration-defined  
 7264 default value for the environment variable *PATH*. Since *PATH* can be modified by the user to  
 7265 include directories that could contain utilities replacing the standard utilities in the Shell and  
 7266 Utilities volume of IEEE Std 1003.1-200x, applications need a way to determine the system-  
 7267 supplied *PATH* environment variable value that contains the correct search path for the standard  
 7268 utilities.

7269 An application could use:

```
7270 confstr(name, (char *)NULL, (size_t)0)
```

7271 to find out how big a buffer is needed for the string value; use *malloc()* to allocate a buffer to  
 7272 hold the string; and call *confstr()* again to get the string. Alternately, it could allocate a fixed,  
 7273 static buffer that is big enough to hold most answers (perhaps 512 or 1 024 bytes), but then use  
 7274 *malloc()* to allocate a larger buffer if it finds that this is too small.

#### 7275 RATIONALE

7276 Application developers can normally determine any configuration variable by means of reading  
 7277 from the stream opened by a call to:

```
7278 popen("command -p getconf variable", "r");
```

7279 The *confstr()* function with a *name* argument of *\_CS\_PATH* returns a string that can be used as a  
 7280 *PATH* environment variable setting that will reference the standard shell and utilities as  
 7281 described in the Shell and Utilities volume of IEEE Std 1003.1-200x.

7282 The *confstr()* function copies the returned string into a buffer supplied by the application instead  
7283 of returning a pointer to a string. This allows a cleaner function in some implementations (such  
7284 as those with lightweight threads) and resolves questions about when the application must copy  
7285 the string returned.

7286 **FUTURE DIRECTIONS**

7287 None.

7288 **SEE ALSO**

7289 *pathconf()*, *sysconf()*, the Base Definitions volume of IEEE Std 1003.1-200x, <**unistd.h**>, the Shell  
7290 and Utilities volume of IEEE Std 1003.1-200x, *c99*

7291 **CHANGE HISTORY**

7292 First released in Issue 4. Derived from the ISO POSIX-2 standard.

7293 **Issue 5**

7294 A table indicating the permissible values of *name* are added to the DESCRIPTION. All those  
7295 marked EX are new in this issue.

7296 **Issue 6**

7297 The Open Group Corrigendum U033/7 is applied. The return value for the case returning the  
7298 size of the buffer now explicitly states that this includes the terminating null.

7299 The following new requirements on POSIX implementations derive from alignment with the  
7300 Single UNIX Specification:

- 7301 • The DESCRIPTION is updated with new arguments which can be used to determine  
7302 configuration strings for C compiler flags, linker/loader flags, and libraries for each different  
7303 supported programming environment. This is a change to support data size neutrality.

7304 The following changes were made to align with the IEEE P1003.1a draft standard:

- 7305 • The DESCRIPTION is updated to include text describing how `_CS_PATH` can be used to  
7306 obtain a *PATH* to access the standard utilities.

7307 The macros associated with the *c89* programming models are marked LEGACY and new  
7308 equivalent macros associated with *c99* are introduced.

7309 **NAME**

7310 conj, conjf, conjl — complex conjugate functions

7311 **SYNOPSIS**

7312 #include &lt;complex.h&gt;

7313 double complex conj(double complex z);

7314 float complex conjf(float complex z);

7315 long double complex conjl(long double complex z);

7316 **DESCRIPTION**7317 cx The functionality described on this reference page is aligned with the ISO C standard. Any  
7318 conflict between the requirements described here and the ISO C standard is unintentional. This  
7319 volume of IEEE Std 1003.1-200x defers to the ISO C standard.7320 These functions shall compute the complex conjugate of *z*, by reversing the sign of its imaginary  
7321 part.7322 **RETURN VALUE**

7323 These functions return the complex conjugate value.

7324 **ERRORS**

7325 No errors are defined.

7326 **EXAMPLES**

7327 None.

7328 **APPLICATION USAGE**

7329 None.

7330 **RATIONALE**

7331 None.

7332 **FUTURE DIRECTIONS**

7333 None.

7334 **SEE ALSO**7335 *carg()*, *cimag()*, *cproj()*, *creal()*, the Base Definitions volume of IEEE Std 1003.1-200x,  
7336 <complex.h>7337 **CHANGE HISTORY**

7338 First released in Issue 6. Derived from the ISO/IEC 9899:1999 standard.

7339 **NAME**

7340 connect — connect a socket

7341 **SYNOPSIS**

7342 #include &lt;sys/socket.h&gt;

7343 int connect(int *socket*, const struct sockaddr \**address*,  
7344 socklen\_t *address\_len*);7345 **DESCRIPTION**7346 The *connect()* function shall attempt to make a connection on a socket. The function takes the  
7347 following arguments:7348 *socket* Specifies the file descriptor associated with the socket.7349 *address* Points to a **sockaddr** structure containing the peer address. The length and  
7350 format of the address depend on the address family of the socket.7351 *address\_len* Specifies the length of the **sockaddr** structure pointed to by the *address*  
7352 argument.7353 If the socket has not already been bound to a local address, *connect()* shall bind it to an address  
7354 which, unless the socket's address family is AF\_UNIX, is an unused local address.7355 If the initiating socket is not connection-mode, then *connect()* shall set the socket's peer address,  
7356 and no connection is made. For SOCK\_DGRAM sockets, the peer address identifies where all  
7357 datagrams are sent on subsequent *send()* functions, and limits the remote sender for subsequent  
7358 *recv()* functions. If *address* is a null address for the protocol, the socket's peer address shall be  
7359 reset.7360 If the initiating socket is connection-mode, then *connect()* shall attempt to establish a connection  
7361 to the address specified by the *address* argument. If the connection cannot be established  
7362 immediately and O\_NONBLOCK is not set for the file descriptor for the socket, *connect()* shall  
7363 block for up to an unspecified timeout interval until the connection is established. If the timeout  
7364 interval expires before the connection is established, *connect()* shall fail and the connection  
7365 attempt shall be aborted. If *connect()* is interrupted by a signal that is caught while blocked  
7366 waiting to establish a connection, *connect()* shall fail and set *errno* to [EINTR], but the connection  
7367 request shall not be aborted, and the connection shall be established asynchronously.7368 If the connection cannot be established immediately and O\_NONBLOCK is set for the file  
7369 descriptor for the socket, *connect()* shall fail and set *errno* to [EINPROGRESS], but the connection  
7370 request shall not be aborted, and the connection shall be established asynchronously.  
7371 Subsequent calls to *connect()* for the same socket, before the connection is established, shall fail  
7372 and set *errno* to [EALREADY].7373 When the connection has been established asynchronously, *select()* and *poll()* shall indicate that  
7374 the file descriptor for the socket is ready for writing.7375 The socket in use may require the process to have appropriate privileges to use the *connect()*  
7376 function.7377 **RETURN VALUE**7378 Upon successful completion, *connect()* shall return 0; otherwise, -1 shall be returned and *errno*  
7379 set to indicate the error.7380 **ERRORS**7381 The *connect()* function shall fail if:

7382 [EADDRNOTAVAIL]

7383 The specified address is not available from the local machine.



7384	[EAFNOSUPPORT]	
7385		The specified address is not a valid address for the address family of the
7386		specified socket.
7387	[EALREADY]	A connection request is already in progress for the specified socket.
7388	[EBADF]	The <i>socket</i> argument is not a valid file descriptor.
7389	[ECONNREFUSED]	
7390		The target address was not listening for connections or refused the connection
7391		request.
7392	[EINPROGRESS]	O_NONBLOCK is set for the file descriptor for the socket and the connection
7393		cannot be immediately established; the connection shall be established
7394		asynchronously.
7395	[EINTR]	The attempt to establish a connection was interrupted by delivery of a signal
7396		that was caught; the connection shall be established asynchronously.
7397	[EISCONN]	The specified socket is connection-mode and is already connected.
7398	[ENETUNREACH]	
7399		No route to the network is present.
7400	[ENOTSOCK]	The <i>socket</i> argument does not refer to a socket.
7401	[EPROTOTYPE]	The specified address has a different type than the socket bound to the
7402		specified peer address.
7403	[ETIMEDOUT]	The attempt to connect timed out before a connection was made.
7404		If the address family of the socket is AF_UNIX, then <i>connect()</i> shall fail if:
7405	[EIO]	An I/O error occurred while reading from or writing to the file system.
7406	[ELOOP]	A loop exists in symbolic links encountered during resolution of the pathname
7407		in <i>address</i> .
7408	[ENAMETOOLONG]	
7409		A component of a pathname exceeded {NAME_MAX} characters, or an entire
7410		pathname exceeded {PATH_MAX} characters.
7411	[ENOENT]	A component of the pathname does not name an existing file or the pathname
7412		is an empty string.
7413	[ENOTDIR]	A component of the path prefix of the pathname in <i>address</i> is not a directory.
7414		The <i>connect()</i> function may fail if:
7415	[EACCES]	Search permission is denied for a component of the path prefix; or write
7416		access to the named socket is denied.
7417	[EADDRINUSE]	Attempt to establish a connection that uses addresses that are already in use.
7418	[ECONNRESET]	Remote host reset the connection request.
7419	[EHOSTUNREACH]	
7420		The destination host cannot be reached (probably because the host is down or
7421		a remote router cannot reach it).
7422	[EINVAL]	The <i>address_len</i> argument is not a valid length for the address family; or
7423		invalid address family in the <b>sockaddr</b> structure.

7424 [ELOOP] More than {SYMLOOP\_MAX} symbolic links were encountered during |  
7425 resolution of the pathname in *address*. |

7426 [ENAMETOOLONG]  
7427 Pathname resolution of a symbolic link produced an intermediate result |  
7428 whose length exceeds {PATH\_MAX}. |

7429 [ENETDOWN] The local network interface used to reach the destination is down.

7430 [ENOBUFS] No buffer space is available.

7431 [EOPNOTSUPP] The socket is listening and cannot be connected.

7432 **EXAMPLES**  
7433 None.

7434 **APPLICATION USAGE**  
7435 If *connect()* fails, the state of the socket is unspecified. Conforming applications should close the |  
7436 file descriptor and create a new socket before attempting to reconnect. |

7437 **RATIONALE**  
7438 None.

7439 **FUTURE DIRECTIONS**  
7440 None.

7441 **SEE ALSO**  
7442 *accept()*, *bind()*, *close()*, *getsockname()*, *poll()*, *select()*, *send()*, *shutdown()*, *socket()*, the Base  
7443 Definitions volume of IEEE Std 1003.1-200x, <**sys/socket.h**>

7444 **CHANGE HISTORY**  
7445 First released in Issue 6. Derived from the XNS, Issue 5.2 specification.  
7446 The wording of the mandatory [ELOOP] error condition is updated, and a second optional  
7447 [ELOOP] error condition is added.

7448 **NAME**

7449 copysign, copysignf, copysignl — number manipulation function

7450 **SYNOPSIS**

7451 #include &lt;math.h&gt;

7452 double copysign(double x, double y);

7453 float copysignf(float x, float y);

7454 long double copysignl(long double x, long double y);

7455 **DESCRIPTION**

7456 **CX** The functionality described on this reference page is aligned with the ISO C standard. Any  
7457 conflict between the requirements described here and the ISO C standard is unintentional. This  
7458 volume of IEEE Std 1003.1-200x defers to the ISO C standard.

7459 These functions shall produce a value with the magnitude of *x* and the sign of *y*. On  
7460 implementations that represent a signed zero but do not treat negative zero consistently in  
7461 arithmetic operations, these functions regard the sign of zero as positive.

7462 **RETURN VALUE**

7463 Upon successful completion, these functions shall return a value with the magnitude of *x* and  
7464 the sign of *y*.

7465 **ERRORS**

7466 No errors are defined.

7467 **EXAMPLES**

7468 None.

7469 **APPLICATION USAGE**

7470 None.

7471 **RATIONALE**

7472 None.

7473 **FUTURE DIRECTIONS**

7474 None.

7475 **SEE ALSO**7476 *signbit()*, the Base Definitions volume of IEEE Std 1003.1-200x, <math.h>7477 **CHANGE HISTORY**

7478 First released in Issue 6. Derived from the ISO/IEC 9899:1999 standard.

7479 **NAME**

7480 cos, cosf, cosl — cosine function

7481 **SYNOPSIS**

7482 #include &lt;math.h&gt;

7483 double cos(double x);

7484 float cosf(float x);

7485 long double cosl(long double x);

7486 **DESCRIPTION**

7487 **CX** The functionality described on this reference page is aligned with the ISO C standard. Any  
 7488 conflict between the requirements described here and the ISO C standard is unintentional. This  
 7489 volume of IEEE Std 1003.1-200x defers to the ISO C standard.

7490 These functions shall compute the cosine of their argument *x*, measured in radians.

7491 An application wishing to check for error situations should set *errno* to zero and call  
 7492 *feclearexcept*(FE\_ALL\_EXCEPT) before calling these functions. On return, if *errno* is non-zero or  
 7493 *fetestexcept*(FE\_INVALID | FE\_DIVBYZERO | FE\_OVERFLOW | FE\_UNDERFLOW) is non-  
 7494 zero, an error has occurred.

7495 **RETURN VALUE**7496 Upon successful completion, these functions shall return the cosine of *x*.7497 **MX** If *x* is NaN, a NaN shall be returned.7498 If *x* is  $\pm 0$ , the value 1.0 shall be returned.

7499 If *x* is  $\pm\text{Inf}$ , a domain error shall occur, and either a NaN (if supported), or an implementation-  
 7500 defined value shall be returned.

7501 **ERRORS**

7502 These functions shall fail if:

7503 **MX** **Domain Error** The *x* argument is  $\pm\text{Inf}$ .

7504 If the integer expression (math\_errhandling & MATH\_ERRNO) is non-zero, |  
 7505 then *errno* shall be set to [EDOM]. If the integer expression (math\_errhandling |  
 7506 & MATH\_ERREXCEPT) is non-zero, then the invalid floating-point exception |  
 7507 shall be raised. |

7508 **EXAMPLES**7509 **Taking the Cosine of a 45-Degree Angle**

7510 #include &lt;math.h&gt;

7511 ...

7512 double radians = 45 \* M\_PI / 180;

7513 double result;

7514 ...

7515 result = cos(radians);

7516 **APPLICATION USAGE**

7517 These functions may lose accuracy when their argument is near an odd multiple of  $\pi/2$  or is far  
 7518 from 0.

7519 On error, the expressions (math\_errhandling & MATH\_ERRNO) and (math\_errhandling &  
 7520 MATH\_ERREXCEPT) are independent of each other, but at least one of them must be non-zero.

7521 **RATIONALE**

7522 None.

7523 **FUTURE DIRECTIONS**

7524 None.

7525 **SEE ALSO**

7526 *acos()*, *feclearexcept()*, *fetetestexcept()*, *isnan()*, *sin()*, *tan()*, the Base Definitions volume of |  
7527 IEEE Std 1003.1-200x, Section 4.18, Treatment of Error Conditions for Mathematical Functions, |  
7528 **<math.h>**

7529 **CHANGE HISTORY**

7530 First released in Issue 1. Derived from Issue 1 of the SVID.

7531 **Issue 5**

7532 The DESCRIPTION is updated to indicate how an application should check for an error. This  
7533 text was previously published in the APPLICATION USAGE section.

7534 **Issue 6**7535 The *cosf()* and *cosl()* functions are added for alignment with the ISO/IEC 9899:1999 standard.

7536 The DESCRIPTION, RETURN VALUE, ERRORS, and APPLICATION USAGE sections are  
7537 revised to align with the ISO/IEC 9899:1999 standard.

7538 IEC 60559:1989 standard floating-point extensions over the ISO/IEC 9899:1999 standard are  
7539 marked.

7540 **NAME**

7541       cosf — cosine function

7542 **SYNOPSIS**

7543       #include &lt;math.h&gt;

7544       float cosf(float x);

7545 **DESCRIPTION**7546       Refer to *cos()*.

7547 **NAME**

7548 cosh, coshf, coshl — hyperbolic cosine functions

7549 **SYNOPSIS**

7550 #include &lt;math.h&gt;

7551 double cosh(double x);

7552 float coshf(float x);

7553 long double coshl(long double x);

7554 **DESCRIPTION**

7555 **CX** The functionality described on this reference page is aligned with the ISO C standard. Any  
 7556 conflict between the requirements described here and the ISO C standard is unintentional. This  
 7557 volume of IEEE Std 1003.1-200x defers to the ISO C standard.

7558 These functions shall compute the hyperbolic cosine of their argument  $x$ .

7559 An application wishing to check for error situations should set *errno* to zero and call  
 7560 *feclearexcept*(FE\_ALL\_EXCEPT) before calling these functions. On return, if *errno* is non-zero or  
 7561 *fetetestexcept*(FE\_INVALID | FE\_DIVBYZERO | FE\_OVERFLOW | FE\_UNDERFLOW) is non-  
 7562 zero, an error has occurred.

7563 **RETURN VALUE**7564 Upon successful completion, these functions shall return the hyperbolic cosine of  $x$ .

7565 If the correct value would cause overflow, a range error shall occur and *cosh*( $x$ ), *coshf*( $x$ ), and  
 7566 *coshl*( $x$ ) shall return the value of the macro HUGE\_VAL, HUGE\_VALF, and HUGE\_VALL,  
 7567 respectively.

7568 **MX** If  $x$  is NaN, a NaN shall be returned.7569 If  $x$  is  $\pm 0$ , the value 1.0 shall be returned.7570 If  $x$  is  $\pm\text{Inf}$ ,  $+\text{Inf}$  shall be returned.7571 **ERRORS**

7572 These functions shall fail if:

7573 **Range Error** The result would cause an overflow.

7574 If the integer expression (math\_errhandling & MATH\_ERRNO) is non-zero, |  
 7575 then *errno* shall be set to [ERANGE]. If the integer expression |  
 7576 (math\_errhandling & MATH\_ERREXCEPT) is non-zero, then the overflow |  
 7577 floating-point exception shall be raised. |

7578 **EXAMPLES**

7579 None.

7580 **APPLICATION USAGE**

7581 On error, the expressions (math\_errhandling & MATH\_ERRNO) and (math\_errhandling &  
 7582 MATH\_ERREXCEPT) are independent of each other, but at least one of them must be non-zero.

7583 For IEEE Std 754-1985 **double**,  $710.5 < |x|$  implies that *cosh*( $x$ ) has overflowed.7584 **RATIONALE**

7585 None.

7586 **FUTURE DIRECTIONS**

7587 None.

7588 **SEE ALSO**

7589 *acosh()*, *feclearexcept()*, *fetestexcept()*, *isnan()*, *sinh()*, *tanh()*, the Base Definitions volume of |  
7590 IEEE Std 1003.1-200x, Section 4.18, Treatment of Error Conditions for Mathematical Functions, |  
7591 **<math.h>**

7592 **CHANGE HISTORY**

7593 First released in Issue 1. Derived from Issue 1 of the SVID.

7594 **Issue 5**

7595 The DESCRIPTION is updated to indicate how an application should check for an error. This  
7596 text was previously published in the APPLICATION USAGE section.

7597 **Issue 6**

7598 The *coshf()* and *coshl()* functions are added for alignment with the ISO/IEC 9899:1999 standard.

7599 The DESCRIPTION, RETURN VALUE, ERRORS, and APPLICATION USAGE sections are  
7600 revised to align with the ISO/IEC 9899:1999 standard.

7601 IEC 60559:1989 standard floating-point extensions over the ISO/IEC 9899:1999 standard are  
7602 marked.



7603 **NAME**

7604       cosl — cosine function

7605 **SYNOPSIS**

7606       #include <math.h>

7607       long double cosl(long double x);

7608 **DESCRIPTION**

7609       Refer to *cos()*.

7610 **NAME**

7611           cpow, cpowf, cpowl — complex power functions

7612 **SYNOPSIS**

7613           #include &lt;complex.h&gt;

7614           double complex cpow(double complex x, double complex y);

7615           float complex cpowf(float complex x, float complex y);

7616           long double complex cpowl(long double complex x,

7617           long double complex y);

7618 **DESCRIPTION**

7619 **CX**       The functionality described on this reference page is aligned with the ISO C standard. Any  
7620 conflict between the requirements described here and the ISO C standard is unintentional. This  
7621 volume of IEEE Std 1003.1-200x defers to the ISO C standard.

7622       These functions shall compute the complex power function  $x^y$ , with a branch cut for the first  
7623 parameter along the negative real axis.

7624 **RETURN VALUE**

7625       These functions shall return the complex power function value.

7626 **ERRORS**

7627       No errors are defined.

7628 **EXAMPLES**

7629       None.

7630 **APPLICATION USAGE**

7631       None.

7632 **RATIONALE**

7633       None.

7634 **FUTURE DIRECTIONS**

7635       None.

7636 **SEE ALSO**7637       *cabs()*, *csqrt()*, the Base Definitions volume of IEEE Std 1003.1-200x, <complex.h>7638 **CHANGE HISTORY**

7639       First released in Issue 6. Derived from the ISO/IEC 9899:1999 standard.

7640 **NAME**

7641 cproj, cprojf, cprojl — complex projection functions

7642 **SYNOPSIS**

7643 #include &lt;complex.h&gt;

7644 double complex cproj(double complex z);

7645 float complex cprojf(float complex z);

7646 long double complex cprojl(long double complex z);

7647 **DESCRIPTION**

7648 cx The functionality described on this reference page is aligned with the ISO C standard. Any  
7649 conflict between the requirements described here and the ISO C standard is unintentional. This  
7650 volume of IEEE Std 1003.1-200x defers to the ISO C standard.

7651 These functions shall compute a projection of  $z$  onto the Riemann sphere:  $z$  projects to  $z$ , except  
7652 that all complex infinities (even those with one infinite part and one NaN part) project to  
7653 positive infinity on the real axis. If  $z$  has an infinite part, then  $cproj(z)$  shall be equivalent to:

7654  $INFINITY + I * copysign(0.0, cimag(z))$ 7655 **RETURN VALUE**

7656 These functions shall return the value of the projection onto the Riemann sphere.

7657 **ERRORS**

7658 No errors are defined.

7659 **EXAMPLES**

7660 None.

7661 **APPLICATION USAGE**

7662 None.

7663 **RATIONALE**

7664 Two topologies are commonly used in complex mathematics: the complex plane with its  
7665 continuum of infinities, and the Riemann sphere with its single infinity. The complex plane is  
7666 better suited for transcendental functions, the Riemann sphere for algebraic functions. The  
7667 complex types with their multiplicity of infinities provide a useful (though imperfect) model for  
7668 the complex plane. The  $cproj()$  function helps model the Riemann sphere by mapping all  
7669 infinities to one, and should be used just before any operation, especially comparisons, that  
7670 might give spurious results for any of the other infinities. Note that a complex value with one  
7671 infinite part and one NaN part is regarded as an infinity, not a NaN, because if one part is  
7672 infinite, the complex value is infinite independent of the value of the other part. For the same  
7673 reason,  $cabs()$  returns an infinity if its argument has an infinite part and a NaN part.

7674 **FUTURE DIRECTIONS**

7675 None.

7676 **SEE ALSO**

7677  $carg()$ ,  $cimag()$ ,  $conj()$ ,  $creal()$ , the Base Definitions volume of IEEE Std 1003.1-200x,  
7678 <complex.h>

7679 **CHANGE HISTORY**

7680 First released in Issue 6. Derived from the ISO/IEC 9899:1999 standard.

7681 **NAME**7682 `creal`, `crealf`, `creall` — complex real functions7683 **SYNOPSIS**7684 `#include <complex.h>`7685 `double creal(double complex z);`7686 `float crealf(float complex z);`7687 `long double creall(long double complex z);`7688 **DESCRIPTION**

7689 **CX** The functionality described on this reference page is aligned with the ISO C standard. Any  
7690 conflict between the requirements described here and the ISO C standard is unintentional. This  
7691 volume of IEEE Std 1003.1-200x defers to the ISO C standard.

7692 These functions shall compute the real part of *z*.7693 **RETURN VALUE**

7694 These functions shall return the real part value.

7695 **ERRORS**

7696 No errors are defined.

7697 **EXAMPLES**

7698 None.

7699 **APPLICATION USAGE**7700 For a variable *z* of complex type:7701 `z == creal(z) + cimag(z)*I`7702 **RATIONALE**

7703 None.

7704 **FUTURE DIRECTIONS**

7705 None.

7706 **SEE ALSO**

7707 `carg()`, `cimag()`, `conj()`, `cproj()`, the Base Definitions volume of IEEE Std 1003.1-200x,  
7708 `<complex.h>`

7709 **CHANGE HISTORY**

7710 First released in Issue 6. Derived from the ISO/IEC 9899:1999 standard.

7711 **NAME**

7712         creat — create a new file or rewrite an existing one

7713 **SYNOPSIS**

7714 OH         #include &lt;sys/stat.h&gt;

7715         #include &lt;fcntl.h&gt;

7716         int creat(const char \*path, mode\_t mode);

7717 **DESCRIPTION**

7718         The function call:

7719         creat(path, mode)

7720         shall be equivalent to:

7721         open(path, O\_WRONLY|O\_CREAT|O\_TRUNC, mode)

7722 **RETURN VALUE**7723         Refer to *open()*.7724 **ERRORS**7725         Refer to *open()*.7726 **EXAMPLES**7727         **Creating a File**7728         The following example creates the file **/tmp/file** with read and write permissions for the file  
7729         owner and read permission for group and others. The resulting file descriptor is assigned to the  
7730         *fd* variable.

7731         #include &lt;fcntl.h&gt;

7732         ...

7733         int fd;

7734         mode\_t mode = S\_IRUSR | S\_IWUSR | S\_IRGRP | S\_IROTH;

7735         char \*filename = "/tmp/file";

7736         ...

7737         fd = creat(filename, mode);

7738         ...

7739 **APPLICATION USAGE**

7740         None.

7741 **RATIONALE**7742         The *creat()* function is redundant. Its services are also provided by the *open()* function. It has  
7743         been included primarily for historical purposes since many existing applications depend on it. It  
7744         is best considered a part of the C binding rather than a function that should be provided in other  
7745         languages.7746 **FUTURE DIRECTIONS**

7747         None.

7748 **SEE ALSO**7749         *open()*, the Base Definitions volume of IEEE Std 1003.1-200x, <fcntl.h>, <sys/stat.h>,  
7750         <sys/types.h>

7751 **CHANGE HISTORY**

7752 First released in Issue 1. Derived from Issue 1 of the SVID.

7753 **Issue 6**

7754 In the SYNOPSIS, the optional include of the `<sys/types.h>` header is removed.

7755 The following new requirements on POSIX implementations derive from alignment with the  
7756 Single UNIX Specification:

- 7757 • The requirement to include `<sys/types.h>` has been removed. Although `<sys/types.h>` was  
7758 required for conforming implementations of previous POSIX specifications, it was not  
7759 required for UNIX applications.

7760 **NAME**7761 crypt — string encoding function (**CRYPT**)7762 **SYNOPSIS**

7763 xSI #include &lt;unistd.h&gt;

7764 char \*crypt(const char \*key, const char \*salt);

7765

7766 **DESCRIPTION**7767 The *crypt()* function is a string encoding function. The algorithm is implementation-defined.7768 The *key* argument points to a string to be encoded. The *salt* argument is a string chosen from the  
7769 set:

7770 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

7771 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

7772 0 1 2 3 4 5 6 7 8 9 . /

7773 The first two characters of this string may be used to perturb the encoding algorithm.

7774 The return value of *crypt()* points to static data that is overwritten by each call.7775 The *crypt()* function need not be reentrant. A function that is not required to be reentrant is not  
7776 required to be thread-safe.7777 **RETURN VALUE**7778 Upon successful completion, *crypt()* shall return a pointer to the encoded string. The first two  
7779 characters of the returned value shall be those of the *salt* argument. Otherwise, it shall return a  
7780 null pointer and set *errno* to indicate the error.7781 **ERRORS**7782 The *crypt()* function shall fail if:

7783 [ENOSYS] The functionality is not supported on this implementation.

7784 **EXAMPLES**7785 **Encoding Passwords**7786 The following example finds a user database entry matching a particular user name and changes  
7787 the current password to a new password. The *crypt()* function generates an encoded version of  
7788 each password. The first call to *crypt()* produces an encoded version of the old password; that  
7789 encoded password is then compared to the password stored in the user database. The second  
7790 call to *crypt()* encodes the new password before it is stored.7791 The *putpwent()* function, used in the following example, is not part of IEEE Std 1003.1-200x.

7792 #include &lt;unistd.h&gt;

7793 #include &lt;pwd.h&gt;

7794 #include &lt;string.h&gt;

7795 #include &lt;stdio.h&gt;

7796 ...

7797 int valid\_change;

7798 int pfd; /\* Integer for file descriptor returned by open(). \*/

7799 FILE \*fpfd; /\* File pointer for use in putpwent(). \*/

7800 struct passwd \*p;

7801 char user[100];

7802 char oldpasswd[100];

7803 char newpasswd[100];

```
7804     char savepasswd[100];
7805     ...
7806     valid_change = 0;
7807     while ((p = getpwent()) != NULL) {
7808         /* Change entry if found. */
7809         if (strcmp(p->pw_name, user) == 0) {
7810             if (strcmp(p->pw_passwd, crypt(oldpasswd, p->pw_passwd)) == 0) {
7811                 strcpy(savepasswd, crypt(newpasswd, user));
7812                 p->pw_passwd = savepasswd;
7813                 valid_change = 1;
7814             }
7815             else {
7816                 fprintf(stderr, "Old password is not valid\n");
7817             }
7818         }
7819         /* Put passwd entry into ptmp. */
7820         putpwent(p, fpfd);
7821     }
```

7822 **APPLICATION USAGE**

7823 The values returned by this function need not be portable among XSI-conformant systems.

7824 **RATIONALE**

7825 None.

7826 **FUTURE DIRECTIONS**

7827 None.

7828 **SEE ALSO**

7829 *encrypt()*, *setkey()*, the Base Definitions volume of IEEE Std 1003.1-200x, <**unistd.h**>

7830 **CHANGE HISTORY**

7831 First released in Issue 1. Derived from Issue 1 of the SVID.

7832 **Issue 5**

7833 Normative text previously in the APPLICATION USAGE section is moved to the  
7834 DESCRIPTION.



7835 **NAME**

7836       csin, csinf, csinl — complex sine functions

7837 **SYNOPSIS**

7838       #include &lt;complex.h&gt;

7839       double complex csin(double complex z);

7840       float complex csinf(float complex z);

7841       long double complex csinl(long double complex z);

7842 **DESCRIPTION**

7843 cx     The functionality described on this reference page is aligned with the ISO C standard. Any  
7844     conflict between the requirements described here and the ISO C standard is unintentional. This  
7845     volume of IEEE Std 1003.1-200x defers to the ISO C standard.

7846       These functions shall compute the complex sine of *z*.7847 **RETURN VALUE**

7848       These functions shall return the complex sine value.

7849 **ERRORS**

7850       No errors are defined.

7851 **EXAMPLES**

7852       None.

7853 **APPLICATION USAGE**

7854       None.

7855 **RATIONALE**

7856       None.

7857 **FUTURE DIRECTIONS**

7858       None.

7859 **SEE ALSO**7860       *casin()*, the Base Definitions volume of IEEE Std 1003.1-200x, <complex.h>7861 **CHANGE HISTORY**

7862       First released in Issue 6. Derived from the ISO/IEC 9899:1999 standard.

7863 **NAME**7864 `csinf` — complex sine functions7865 **SYNOPSIS**7866 `#include <complex.h>`7867 `float complex csinf(float complex z);`7868 **DESCRIPTION**7869 Refer to *csin()*.

7870 **NAME**

7871       csinh, csinhf, csinhl — complex hyperbolic sine functions

7872 **SYNOPSIS**

7873       #include &lt;complex.h&gt;

7874       double complex csinh(double complex *z*);7875       float complex csinhf(float complex *z*);7876       long double complex csinhl(long double complex *z*);7877 **DESCRIPTION**

7878 **CX**     The functionality described on this reference page is aligned with the ISO C standard. Any  
7879     conflict between the requirements described here and the ISO C standard is unintentional. This  
7880     volume of IEEE Std 1003.1-200x defers to the ISO C standard.

7881       These functions shall compute the complex hyperbolic sine of *z*.7882 **RETURN VALUE**

7883       These functions shall return the complex hyperbolic sine value.

7884 **ERRORS**

7885       No errors are defined.

7886 **EXAMPLES**

7887       None.

7888 **APPLICATION USAGE**

7889       None.

7890 **RATIONALE**

7891       None.

7892 **FUTURE DIRECTIONS**

7893       None.

7894 **SEE ALSO**7895       *casinh()*, the Base Definitions volume of IEEE Std 1003.1-200x, <complex.h>7896 **CHANGE HISTORY**

7897       First released in Issue 6. Derived from the ISO/IEC 9899:1999 standard.

7898 **NAME**

7899       csinl — complex sine functions

7900 **SYNOPSIS**

7901       #include &lt;complex.h&gt;

7902       long double complex csinl(long double complex z);

7903 **DESCRIPTION**7904       Refer to *csin()*.

7905 **NAME**

7906 csqrt, csqrtf, csqrtl — complex square root functions

7907 **SYNOPSIS**

7908 #include &lt;complex.h&gt;

7909 double complex csqrt(double complex z);

7910 float complex csqrtf(float complex z);

7911 long double complex csqrtl(long double complex z);

7912 **DESCRIPTION**

7913 cx The functionality described on this reference page is aligned with the ISO C standard. Any  
7914 conflict between the requirements described here and the ISO C standard is unintentional. This  
7915 volume of IEEE Std 1003.1-200x defers to the ISO C standard.

7916 These functions shall compute the complex square root of  $z$ , with a branch cut along the  
7917 negative real axis.

7918 **RETURN VALUE**

7919 These functions shall return the complex square root value, in the range of the right half-plane  
7920 (including the imaginary axis).

7921 **ERRORS**

7922 No errors are defined.

7923 **EXAMPLES**

7924 None.

7925 **APPLICATION USAGE**

7926 None.

7927 **RATIONALE**

7928 None.

7929 **FUTURE DIRECTIONS**

7930 None.

7931 **SEE ALSO**7932 *cabs()*, *cpow()*, the Base Definitions volume of IEEE Std 1003.1-200x, <complex.h>7933 **CHANGE HISTORY**

7934 First released in Issue 6. Derived from the ISO/IEC 9899:1999 standard.

7935 **NAME**

7936           ctan, ctanf, ctanl — complex tangent functions

7937 **SYNOPSIS**

7938           #include &lt;complex.h&gt;

7939           double complex ctan(double complex *z*);7940           float complex ctanf(float complex *z*);7941           long double complex ctanl(long double complex *z*);7942 **DESCRIPTION**7943 **CX**       The functionality described on this reference page is aligned with the ISO C standard. Any  
7944 conflict between the requirements described here and the ISO C standard is unintentional. This  
7945 volume of IEEE Std 1003.1-200x defers to the ISO C standard.7946           These functions shall compute the complex tangent of *z*.7947 **RETURN VALUE**

7948           These functions shall return the complex tangent value.

7949 **ERRORS**

7950           No errors are defined.

7951 **EXAMPLES**

7952           None.

7953 **APPLICATION USAGE**

7954           None.

7955 **RATIONALE**

7956           None.

7957 **FUTURE DIRECTIONS**

7958           None.

7959 **SEE ALSO**7960           *catan()*, the Base Definitions volume of IEEE Std 1003.1-200x, <complex.h>7961 **CHANGE HISTORY**

7962           First released in Issue 6. Derived from the ISO/IEC 9899:1999 standard.

7963 **NAME**

7964       ctanf — complex tangent functions

7965 **SYNOPSIS**

7966       #include &lt;complex.h&gt;

7967       float complex ctanf(float complex z);

7968 **DESCRIPTION**7969       Refer to *ctan()*.

7970 **NAME**

7971 ctanh, ctanhf, ctanhl — complex hyperbolic tangent functions

7972 **SYNOPSIS**

7973 #include &lt;complex.h&gt;

7974 double complex ctanh(double complex z);

7975 float complex ctanhf(float complex z);

7976 long double complex ctanhl(long double complex z);

7977 **DESCRIPTION**7978 **CX** The functionality described on this reference page is aligned with the ISO C standard. Any  
7979 conflict between the requirements described here and the ISO C standard is unintentional. This  
7980 volume of IEEE Std 1003.1-200x defers to the ISO C standard.7981 These functions shall compute the complex hyperbolic tangent of  $z$ .7982 **RETURN VALUE**

7983 These functions shall return the complex hyperbolic tangent value.

7984 **ERRORS**

7985 No errors are defined.

7986 **EXAMPLES**

7987 None.

7988 **APPLICATION USAGE**

7989 None.

7990 **RATIONALE**

7991 None.

7992 **FUTURE DIRECTIONS**

7993 None.

7994 **SEE ALSO**7995 *catanh()*, the Base Definitions volume of IEEE Std 1003.1-200x, <complex.h>7996 **CHANGE HISTORY**

7997 First released in Issue 6. Derived from the ISO/IEC 9899:1999 standard.



7998 **NAME**

7999       ctanl — complex tangent functions

8000 **SYNOPSIS**

8001       #include <complex.h>

8002       long double complex ctanl(long double complex z);

8003 **DESCRIPTION**

8004       Refer to *ctan()*.

8005 **NAME**

8006 ctermid — generate a pathname for controlling terminal |

8007 **SYNOPSIS**

8008 cx #include &lt;stdio.h&gt; |

8009 char \*ctermid(char \*s); |

8010 |

8011 **DESCRIPTION**

8012 The *ctermid()* function shall generate a string that, when used as a pathname, refers to the |  
 8013 current controlling terminal for the current process. If *ctermid()* returns a pathname, access to the |  
 8014 file is not guaranteed. |

8015 If the application uses any of the `_POSIX_THREAD_SAFE_FUNCTIONS` or `_POSIX_THREADS` |  
 8016 functions, it shall ensure that the *ctermid()* function is called with a non-NULL parameter. |

8017 **RETURN VALUE**

8018 If *s* is a null pointer, the string shall be generated in an area that may be static (and therefore may |  
 8019 be overwritten by each call), the address of which shall be returned. Otherwise, *s* is assumed to |  
 8020 point to a character array of at least `L_ctermid` bytes; the string is placed in this array and the |  
 8021 value of *s* shall be returned. The symbolic constant `L_ctermid` is defined in `<stdio.h>`, and shall |  
 8022 have a value greater than 0. |

8023 The *ctermid()* function shall return an empty string if the pathname that would refer to the |  
 8024 controlling terminal cannot be determined, or if the function is unsuccessful. |

8025 **ERRORS**

8026 No errors are defined.

8027 **EXAMPLES**8028 **Determining the Controlling Terminal for the Current Process**

8029 The following example returns a pointer to a string that identifies the controlling terminal for the |  
 8030 current process. The pathname for the terminal is stored in the array pointed to by the *ptr* |  
 8031 argument, which has a size of `L_ctermid` bytes, as indicated by the *term* argument. |

8032 #include &lt;stdio.h&gt;

8033 ...

8034 char term[L\_ctermid];

8035 char \*ptr;

8036 ptr = ctermid(term);

8037 **APPLICATION USAGE**

8038 The difference between *ctermid()* and *ttyname()* is that *ttyname()* must be handed a file |  
 8039 descriptor and return a path of the terminal associated with that file descriptor, while *ctermid()* |  
 8040 returns a string (such as `"/dev/tty"`) that refers to the current controlling terminal if used as a |  
 8041 pathname. |

8042 **RATIONALE**

8043 `L_ctermid` must be defined appropriately for a given implementation and must be greater than |  
 8044 zero so that array declarations using it are accepted by the compiler. The value includes the |  
 8045 terminating null byte. |

8046 Conforming applications that use threads cannot call *ctermid()* with NULL as the parameter if |  
 8047 either `_POSIX_THREAD_SAFE_FUNCTIONS` or `_POSIX_THREADS` is defined. If *s* is not |  
 8048 NULL, the *ctermid()* function generates a string that, when used as a pathname, refers to the |

8049 current controlling terminal for the current process. If *s* is NULL, the return value of *ctermid()* is  
8050 undefined.

8051 There is no additional burden on the programmer—changing to use a hypothetical thread-safe  
8052 version of *ctermid()* along with allocating a buffer is more of a burden than merely allocating a  
8053 buffer. Application code should not assume that the returned string is short, as some  
8054 implementations have more than two pathname components before reaching a logical device  
8055 name. |

8056 **FUTURE DIRECTIONS**

8057 None.

8058 **SEE ALSO**

8059 *ttyname()*, the Base Definitions volume of IEEE Std 1003.1-200x, <**stdio.h**>

8060 **CHANGE HISTORY**

8061 First released in Issue 1. Derived from Issue 1 of the SVID.

8062 **Issue 5**

8063 The DESCRIPTION is updated for alignment with the POSIX Threads Extension.

8064 **Issue 6**

8065 The DESCRIPTION is updated to avoid use of the term “must” for application requirements.

8066 **NAME**

8067 ctime, ctime\_r — convert a time value to date and time string

8068 **SYNOPSIS**

8069 #include &lt;time.h&gt;

8070 char \*ctime(const time\_t \*clock);

8071 TSF char \*ctime\_r(const time\_t \*clock, char \*buf);

8072

8073 **DESCRIPTION**

8074 CX For *ctime()*: The functionality described on this reference page is aligned with the ISO C |  
 8075 standard. Any conflict between the requirements described here and the ISO C standard is |  
 8076 unintentional. This volume of IEEE Std 1003.1-200x defers to the ISO C standard.

8077 The *ctime()* function shall convert the time pointed to by *clock*, representing time in seconds |  
 8078 since the Epoch, to local time in the form of a string. It shall be equivalent to: |

8079 asctime(localtime(clock))

8080 CX The *asctime()*, *ctime()*, *gmtime()*, and *localtime()* functions shall return values in one of two static |  
 8081 objects: a broken-down time structure and an array of **char**. Execution of any of the functions |  
 8082 may overwrite the information returned in either of these objects by any of the other functions.

8083 The *ctime()* function need not be reentrant. A function that is not required to be reentrant is not |  
 8084 required to be thread-safe.

8085 TSF The *ctime\_r()* function shall convert the calendar time pointed to by *clock* to local time in exactly |  
 8086 the same form as *ctime()* and puts the string into the array pointed to by *buf* (which shall be at |  
 8087 least 26 bytes in size) and return *buf*.

8088 Unlike *ctime()*, the thread-safe version *ctime\_r()* is not required to set *tzname*.

8089 **RETURN VALUE**

8090 The *ctime()* function shall return the pointer returned by *asctime()* with that broken-down time |  
 8091 as an argument.

8092 TSF Upon successful completion, *ctime\_r()* shall return a pointer to the string pointed to by *buf*. |  
 8093 When an error is encountered, a null pointer shall be returned.

8094 **ERRORS**

8095 No errors are defined.

8096 **EXAMPLES**

8097 None.

8098 **APPLICATION USAGE**

8099 Values for the broken-down time structure can be obtained by calling *gmtime()* or *localtime()*.  
 8100 The *ctime()* function is included for compatibility with older implementations, and does not  
 8101 support localized date and time formats. Applications should use the *strptime()* function to  
 8102 achieve maximum portability.

8103 The *ctime\_r()* function is thread-safe and shall return values in a user-supplied buffer instead of  
 8104 possibly using a static data area that may be overwritten by each call.

8105 **RATIONALE**

8106 None.

8107 **FUTURE DIRECTIONS**

8108 None.

8109 **SEE ALSO**8110 *asctime()*, *clock()*, *difftime()*, *gmtime()*, *localtime()*, *mktime()*, *strftime()*, *strptime()*, *time()*, *utime()*,  
8111 the Base Definitions volume of IEEE Std 1003.1-200x, <**time.h**>8112 **CHANGE HISTORY**

8113 First released in Issue 1. Derived from Issue 1 of the SVID.

8114 **Issue 5**8115 Normative text previously in the APPLICATION USAGE section is moved to the  
8116 DESCRIPTION.8117 The *ctime\_r()* function is included for alignment with the POSIX Threads Extension.8118 A note indicating that the *ctime()* function need not be reentrant is added to the DESCRIPTION.8119 **Issue 6**

8120 Extensions beyond the ISO C standard are now marked.

8121 In the DESCRIPTION, the note about reentrancy is expanded to cover thread-safety.

8122 The APPLICATION USAGE section is updated to include a note on the thread-safe function and  
8123 its avoidance of possibly using a static data area.

8124 **NAME**

8125 daylight — daylight savings time flag

8126 **SYNOPSIS**

8127 xSI #include <time.h>

8128 extern int daylight;

8129

8130 **DESCRIPTION**

8131 Refer to *tzset()*.

## 8132 NAME

8133 dbm\_clearerr, dbm\_close, dbm\_delete, dbm\_error, dbm\_fetch, dbm\_firstkey, dbm\_nextkey,  
8134 dbm\_open, dbm\_store — database functions

## 8135 SYNOPSIS

```
8136 xSI #include <ndbm.h>
8137
8137 int dbm_clearerr(DBM *db);
8138 void dbm_close(DBM *db);
8139 int dbm_delete(DBM *db, datum key);
8140 int dbm_error(DBM *db);
8141 datum dbm_fetch(DBM *db, datum key);
8142 datum dbm_firstkey(DBM *db);
8143 datum dbm_nextkey(DBM *db);
8144 DBM *dbm_open(const char *file, int open_flags, mode_t file_mode);
8145 int dbm_store(DBM *db, datum key, datum content, int store_mode);
8146
```

## 8147 DESCRIPTION

8148 These functions create, access, and modify a database.

8149 A **datum** consists of at least two members, *dptr* and *dsize*. The *dptr* member points to an object  
8150 that is *dsize* bytes in length. Arbitrary binary data, as well as character strings, may be stored in  
8151 the object pointed to by *dptr*.

8152 The database is stored in two files. One file is a directory containing a bit map of keys and has  
8153 **.dir** as its suffix. The second file contains all data and has **.pag** as its suffix.

8154 The *dbm\_open()* function shall open a database. The *file* argument to the function is the  
8155 pathname of the database. The function opens two files named *file.dir* and *file.pag*. The  
8156 *open\_flags* argument has the same meaning as the *flags* argument of *open()* except that a database  
8157 opened for write-only access opens the files for read and write access and the behavior of the  
8158 O\_APPEND flag is unspecified. The *file\_mode* argument has the same meaning as the third  
8159 argument of *open()*.

8160 The *dbm\_close()* function shall close a database. The application shall ensure that argument *db* is  
8161 a pointer to a **dbm** structure that has been returned from a call to *dbm\_open()*.

8162 These database functions shall support an internal block size large enough to support  
8163 key/content pairs of at least 1023 bytes.

8164 The *dbm\_fetch()* function shall read a record from a database. The argument *db* is a pointer to a  
8165 database structure that has been returned from a call to *dbm\_open()*. The argument *key* is a  
8166 **datum** that has been initialized by the application to the value of the key that matches the key of  
8167 the record the program is fetching.

8168 The *dbm\_store()* function shall write a record to a database. The argument *db* is a pointer to a  
8169 database structure that has been returned from a call to *dbm\_open()*. The argument *key* is a  
8170 **datum** that has been initialized by the application to the value of the key that identifies (for  
8171 subsequent reading, writing, or deleting) the record the application is writing. The argument  
8172 *content* is a **datum** that has been initialized by the application to the value of the record the  
8173 program is writing. The argument *store\_mode* controls whether *dbm\_store()* replaces any pre-  
8174 existing record that has the same key that is specified by the *key* argument. The application shall  
8175 set *store\_mode* to either DBM\_INSERT or DBM\_REPLACE. If the database contains a record that  
8176 matches the *key* argument and *store\_mode* is DBM\_REPLACE, the existing record shall be  
8177 replaced with the new record. If the database contains a record that matches the *key* argument  
8178 and *store\_mode* is DBM\_INSERT, the existing record shall be left unchanged and the new record

8179 ignored. If the database does not contain a record that matches the *key* argument and *store\_mode* |  
8180 is either DBM\_INSERT or DBM\_REPLACE, the new record shall be inserted in the database. |

8181 If the sum of a key/content pair exceeds the internal block size, the result is unspecified. |  
8182 Moreover, the application shall ensure that all key/content pairs that hash together fit on a |  
8183 single block. The *dbm\_store()* function shall return an error in the event that a disk block fills |  
8184 with inseparable data.

8185 The *dbm\_delete()* function shall delete a record and its key from the database. The argument *db* is  
8186 a pointer to a database structure that has been returned from a call to *dbm\_open()*. The argument  
8187 *key* is a **datum** that has been initialized by the application to the value of the key that identifies  
8188 the record the program is deleting.

8189 The *dbm\_firstkey()* function shall return the first key in the database. The argument *db* is a  
8190 pointer to a database structure that has been returned from a call to *dbm\_open()*.

8191 The *dbm\_nextkey()* function shall return the next key in the database. The argument *db* is a  
8192 pointer to a database structure that has been returned from a call to *dbm\_open()*. The application  
8193 shall ensure that the *dbm\_firstkey()* function is called before calling *dbm\_nextkey()*. Subsequent  
8194 calls to *dbm\_nextkey()* return the next key until all of the keys in the database have been  
8195 returned.

8196 The *dbm\_error()* function shall return the error condition of the database. The argument *db* is a  
8197 pointer to a database structure that has been returned from a call to *dbm\_open()*.

8198 The *dbm\_clearerr()* function shall clear the error condition of the database. The argument *db* is a  
8199 pointer to a database structure that has been returned from a call to *dbm\_open()*.

8200 The *dptr* pointers returned by these functions may point into static storage that may be changed |  
8201 by subsequent calls.

8202 These functions need not be reentrant. A function that is not required to be reentrant is not  
8203 required to be thread-safe.

8204 **RETURN VALUE**

8205 The *dbm\_store()* and *dbm\_delete()* functions shall return 0 when they succeed and a negative  
8206 value when they fail.

8207 The *dbm\_store()* function shall return 1 if it is called with a *flags* value of DBM\_INSERT and the  
8208 function finds an existing record with the same key.

8209 The *dbm\_error()* function shall return 0 if the error condition is not set and return a non-zero  
8210 value if the error condition is set.

8211 The return value of *dbm\_clearerr()* is unspecified.

8212 The *dbm\_firstkey()* and *dbm\_nextkey()* functions shall return a key **datum**. When the end of the  
8213 database is reached, the *dptr* member of the key is a null pointer. If an error is detected, the *dptr*  
8214 member of the key shall be a null pointer and the error condition of the database shall be set.

8215 The *dbm\_fetch()* function shall return a content **datum**. If no record in the database matches the  
8216 key or if an error condition has been detected in the database, the *dptr* member of the content  
8217 shall be a null pointer.

8218 The *dbm\_open()* function shall return a pointer to a database structure. If an error is detected  
8219 during the operation, *dbm\_open()* shall return a **(DBM \*)0**.



**8220 ERRORS**

8221 No errors are defined.

**8222 EXAMPLES**

8223 None.

**8224 APPLICATION USAGE**

8225 The following code can be used to traverse the database:

```
8226 for(key = dbm_firstkey(db); key.dptr != NULL; key = dbm_nextkey(db))
```

8227 The *dbm\_* functions provided in this library should not be confused in any way with those of a  
8228 general-purpose database management system. These functions do not provide for multiple  
8229 search keys per entry, they do not protect against multi-user access (in other words they do not  
8230 lock records or files), and they do not provide the many other useful database functions that are  
8231 found in more robust database management systems. Creating and updating databases by use of  
8232 these functions is relatively slow because of data copies that occur upon hash collisions. These  
8233 functions are useful for applications requiring fast lookup of relatively static information that is  
8234 to be indexed by a single key.

8235 The *dbm\_delete()* function need not physically reclaim file space, although it does make it  
8236 available for reuse by the database.

8237 After calling *dbm\_store()* or *dbm\_delete()* during a pass through the keys by *dbm\_firstkey()* and  
8238 *dbm\_nextkey()*, the application should reset the database by calling *dbm\_firstkey()* before again  
8239 calling *dbm\_nextkey()*. The contents of these files are unspecified and may not be portable.

**8240 RATIONALE**

8241 None.

**8242 FUTURE DIRECTIONS**

8243 None.

**8244 SEE ALSO**

8245 *open()*, the Base Definitions volume of IEEE Std 1003.1-200x, <**ndbm.h**>

**8246 CHANGE HISTORY**

8247 First released in Issue 4, Version 2.

**8248 Issue 5**

8249 Moved from X/OPEN UNIX extension to BASE.

8250 Normative text previously in the APPLICATION USAGE section is moved to the  
8251 DESCRIPTION.

8252 A note indicating that these functions need not be reentrant is added to the DESCRIPTION.

**8253 Issue 6**

8254 The DESCRIPTION is updated to avoid use of the term “must” for application requirements.

8255 **NAME**

8256 difftime — compute the difference between two calendar time values

8257 **SYNOPSIS**

8258 #include <time.h>

8259 double difftime(time\_t *time1*, time\_t *time0*);

8260 **DESCRIPTION**

8261 cx The functionality described on this reference page is aligned with the ISO C standard. Any  
8262 conflict between the requirements described here and the ISO C standard is unintentional. This  
8263 volume of IEEE Std 1003.1-200x defers to the ISO C standard.

8264 The *difftime()* function shall compute the difference between two calendar times (as returned by  
8265 *time()*): *time1*– *time0*.

8266 **RETURN VALUE**

8267 The *difftime()* function shall return the difference expressed in seconds as a type **double**.

8268 **ERRORS**

8269 No errors are defined.

8270 **EXAMPLES**

8271 None.

8272 **APPLICATION USAGE**

8273 None.

8274 **RATIONALE**

8275 None.

8276 **FUTURE DIRECTIONS**

8277 None.

8278 **SEE ALSO**

8279 *asctime()*, *clock()*, *ctime()*, *gmtime()*, *localtime()*, *mktime()*, *strftime()*, *strptime()*, *time()*, *utime()*,  
8280 the Base Definitions volume of IEEE Std 1003.1-200x, <**time.h**>

8281 **CHANGE HISTORY**

8282 First released in Issue 4. Derived from the ISO C standard.

8283 **NAME**

8284            dirname — report the parent directory name of a file pathname |

8285 **SYNOPSIS**

8286 xSI        #include &lt;libgen.h&gt;

8287            char \*dirname(char \*path);

8288

8289 **DESCRIPTION**8290            The *dirname()* function shall take a pointer to a character string that contains a pathname, and |  
8291            return a pointer to a string that is a pathname of the parent directory of that file. Trailing '/' |  
8292            characters in the path are not counted as part of the path.8293            If *path* does not contain a '/', then *dirname()* shall return a pointer to the string ".". If *path* is a  
8294            null pointer or points to an empty string, *dirname()* shall return a pointer to the string ".".8295            The *dirname()* function need not be reentrant. A function that is not required to be reentrant is  
8296            not required to be thread-safe.8297 **RETURN VALUE**8298            The *dirname()* function shall return a pointer to a string that is the parent directory of *path*. If  
8299            *path* is a null pointer or points to an empty string, a pointer to a string "." is returned.8300            The *dirname()* function may modify the string pointed to by *path*, and may return a pointer to  
8301            static storage that may then be overwritten by subsequent calls to *dirname()*.8302 **ERRORS**

8303            No errors are defined.

8304 **EXAMPLES**8305            The following code fragment reads a pathname, changes the current working directory to the |  
8306            parent directory, and opens the file.8307            char path[MAXPATHLEN], \*pathcopy;  
8308            int fd;  
8309            fgets(path, MAXPATHLEN, stdin);  
8310            pathcopy = strdup(path);  
8311            chdir(dirname(pathcopy));  
8312            fd = open(basename(path), O\_RDONLY);8313 **Sample Input and Output Strings for dirname()**8314            In the following table, the input string is the value pointed to by *path*, and the output string is  
8315            the return value of the *dirname()* function.

Input String	Output String
"/usr/lib"	"/usr"
"/usr/"	"/"
"usr"	."
"/"	"/"
."	."
".."	."

8316

8317

8318

8319

8320

8321

8322

8323 **Changing the Current Directory to the Parent Directory**

8324 The following program fragment reads a pathname, changes the current working directory to |  
8325 the parent directory, and opens the file.

```
8326 #include <unistd.h>
8327 #include <limits.h>
8328 #include <stdio.h>
8329 #include <fcntl.h>
8330 #include <string.h>
8331 #include <libgen.h>
8332 ...
8333 char path[PATH_MAX], *pathcopy;
8334 int fd;
8335 ...
8336 fgets(path, PATH_MAX, stdin);
8337 pathcopy = strdup(path);
8338 chdir(dirname(pathcopy));
8339 fd = open(basename(path), O_RDONLY);
```

8340 **APPLICATION USAGE**

8341 The *dirname()* and *basename()* functions together yield a complete pathname. The expression |  
8342 *dirname(path)* obtains the pathname of the directory where *basename(path)* is found. |

8343 Since the meaning of the leading `"/"` is implementation-defined, *dirname("/foo)* may return  
8344 either `"/"` or `'/'` (but nothing else).

8345 **RATIONALE**

8346 None.

8347 **FUTURE DIRECTIONS**

8348 None.

8349 **SEE ALSO**

8350 *basename()*, the Base Definitions volume of IEEE Std 1003.1-200x, `<libgen.h>`

8351 **CHANGE HISTORY**

8352 First released in Issue 4, Version 2.

8353 **Issue 5**

8354 Moved from X/OPEN UNIX extension to BASE.

8355 Normative text previously in the APPLICATION USAGE section is moved to the  
8356 DESCRIPTION.

8357 A note indicating that this function need not be reentrant is added to the DESCRIPTION.

8358 **NAME**8359 `div` — compute the quotient and remainder of an integer division8360 **SYNOPSIS**8361 `#include <stdlib.h>`8362 `div_t div(int numer, int denom);`8363 **DESCRIPTION**

8364 `CX` The functionality described on this reference page is aligned with the ISO C standard. Any  
8365 conflict between the requirements described here and the ISO C standard is unintentional. This  
8366 volume of IEEE Std 1003.1-200x defers to the ISO C standard.

8367 The `div()` function shall compute the quotient and remainder of the division of the numerator  
8368 `numer` by the denominator `denom`. If the division is inexact, the resulting quotient is the integer  
8369 of lesser magnitude that is the nearest to the algebraic quotient. If the result cannot be  
8370 represented, the behavior is undefined; otherwise, `quot*denom+rem` shall equal `numer`.

8371 **RETURN VALUE**

8372 The `div()` function shall return a structure of type `div_t`, comprising both the quotient and the  
8373 remainder. The structure includes the following members, in any order:

8374 `int quot; /* quotient */`8375 `int rem; /* remainder */`8376 **ERRORS**

8377 No errors are defined.

8378 **EXAMPLES**

8379 None.

8380 **APPLICATION USAGE**

8381 None.

8382 **RATIONALE**

8383 None.

8384 **FUTURE DIRECTIONS**

8385 None.

8386 **SEE ALSO**8387 `ldiv()`, the Base Definitions volume of IEEE Std 1003.1-200x, `<stdlib.h>`8388 **CHANGE HISTORY**

8389 First released in Issue 4. Derived from the ISO C standard.

8390 **NAME**8391           dldclose — close a *dlopen()* object8392 **SYNOPSIS**

8393 XSI       #include &lt;dldfcn.h&gt;

8394       int dldclose(void \*handle);

8395

8396 **DESCRIPTION**8397       The *dldclose()* function shall inform the system that the object referenced by a *handle* returned |  
8398       from a previous *dlopen()* invocation is no longer needed by the application.8399       The use of *dldclose()* reflects a statement of intent on the part of the process, but does not create  
8400       any requirement upon the implementation, such as removal of the code or symbols referenced  
8401       by *handle*. Once an object has been closed using *dldclose()* an application should assume that its  
8402       symbols are no longer available to *dlsym()*. All objects loaded automatically as a result of  
8403       invoking *dlopen()* on the referenced object shall also be closed if this is the last reference to it. |8404       Although a *dldclose()* operation is not required to remove structures from an address space,  
8405       neither is an implementation prohibited from doing so. The only restriction on such a removal is  
8406       that no object shall be removed to which references have been relocated, until or unless all such  
8407       references are removed. For instance, an object that had been loaded with a *dlopen()* operation  
8408       specifying the `RTLD_GLOBAL` flag might provide a target for dynamic relocations performed in  
8409       the processing of other objects—in such environments, an application may assume that no  
8410       relocation, once made, shall be undone or remade unless the object requiring the relocation has  
8411       itself been removed.8412 **RETURN VALUE**8413       If the referenced object was successfully closed, *dldclose()* shall return 0. If the object could not be  
8414       closed, or if *handle* does not refer to an open object, *dldclose()* shall return a non-zero value. More  
8415       detailed diagnostic information shall be available through *dlderror()*.8416 **ERRORS**

8417       No errors are defined.

8418 **EXAMPLES**8419       The following example illustrates use of *dlopen()* and *dldclose()*:8420       ...  
8421       /\* Open a dynamic library and then close it ... \*/  
8422       #include <dldfcn.h>  
8423       void \*mylib;  
8424       int eret;  
8425       mylib = dlopen("mylib.so", RTLD\_LOCAL | RTLD\_LAZY); |  
8426       ... |  
8427       eret = dldclose(mylib);  
8428       ...8429 **APPLICATION USAGE**8430       A conforming application should employ a *handle* returned from a *dlopen()* invocation only |  
8431       within a given scope bracketed by the *dlopen()* and *dldclose()* operations. Implementations are  
8432       free to use reference counting or other techniques such that multiple calls to *dlopen()* referencing  
8433       the same object may return the same object for *handle*. Implementations are also free to reuse a  
8434       *handle*. For these reasons, the value of a *handle* must be treated as an opaque object by the  
8435       application, used only in calls to *dlsym()* and *dldclose()*.

8436 **RATIONALE**

8437       None.

8438 **FUTURE DIRECTIONS**

8439       None.

8440 **SEE ALSO**8441       *derror()*, *dlopen()*, *dlsym()*, the Base Definitions volume of IEEE Std 1003.1-200x, <dlfcn.h>8442 **CHANGE HISTORY**

8443       First released in Issue 5.

8444 **Issue 6**8445       The DESCRIPTION is updated to say that the referenced object is closed “if this is the last  
8446       reference to it”.

8447 **NAME**8448 `dlderror` — get diagnostic information8449 **SYNOPSIS**8450 XSI `#include <dldfcn.h>`8451 `char *dlderror(void);`

8452

8453 **DESCRIPTION**

8454 The `dlderror()` function shall return a null-terminated character string (with no trailing <newline>)  
8455 that describes the last error that occurred during dynamic linking processing. If no dynamic  
8456 linking errors have occurred since the last invocation of `dlderror()`, `dlderror()` shall return NULL.  
8457 Thus, invoking `dlderror()` a second time, immediately following a prior invocation, shall result in  
8458 NULL being returned.

8459 The `dlderror()` function need not be reentrant. A function that is not required to be reentrant is not  
8460 required to be thread-safe.

8461 **RETURN VALUE**

8462 If successful, `dlderror()` shall return a null-terminated character string; otherwise, NULL shall be  
8463 returned.

8464 **ERRORS**

8465 No errors are defined.

8466 **EXAMPLES**

8467 The following example prints out the last dynamic linking error:

```
8468 ...  
8469 #include <dldfcn.h>  
8470 char *errstr;  
8471 errstr = dlderror();  
8472 if (errstr != NULL)  
8473 printf ("A dynamic linking error occurred: (%s)\n", errstr);  
8474 ...
```

8475 **APPLICATION USAGE**

8476 The messages returned by `dlderror()` may reside in a static buffer that is overwritten on each call  
8477 to `dlderror()`. Application code should not write to this buffer. Programs wishing to preserve an  
8478 error message should make their own copies of that message. Depending on the application  
8479 environment with respect to asynchronous execution events, such as signals or other  
8480 asynchronous computation sharing the address space, conforming applications should use a  
8481 critical section to retrieve the error pointer and buffer. |

8482 **RATIONALE**

8483 None.

8484 **FUTURE DIRECTIONS**

8485 None.

8486 **SEE ALSO**8487 `dldclose()`, `dldopen()`, `dldsym()`, the Base Definitions volume of IEEE Std 1003.1-200x, <dldfcn.h>



8488 **CHANGE HISTORY**

8489           First released in Issue 5.

8490 **Issue 6**

8491           In the DESCRIPTION the note about reentrancy and thread-safety is added.

## 8492 NAME

8493 dlopen — gain access to an executable object file

## 8494 SYNOPSIS

8495 xSI #include &lt;dlfcn.h&gt;

8496 void \*dlopen(const char \*file, int mode);

8497

## 8498 DESCRIPTION

8499 The *dlopen()* function shall make an executable object file specified by *file* available to the calling  
 8500 program. The class of files eligible for this operation and the manner of their construction are  
 8501 implementation-defined, though typically such files are executable objects such as shared  
 8502 libraries, relocatable files, or programs. Note that some implementations permit the construction  
 8503 of dependencies between such objects that are embedded within files. In such cases, a *dlopen()*  
 8504 operation shall load such dependencies in addition to the object referenced by *file*.  
 8505 Implementations may also impose specific constraints on the construction of programs that can  
 8506 employ *dlopen()* and its related services.

8507 A successful *dlopen()* shall return a *handle* which the caller may use on subsequent calls to  
 8508 *dlsym()* and *dlclose()*. The value of this *handle* should not be interpreted in any way by the caller.

8509 The *file* argument is used to construct a pathname to the object file. If *file* contains a slash  
 8510 character, the *file* argument is used as the pathname for the file. Otherwise, *file* is used in an  
 8511 implementation-defined manner to yield a pathname.

8512 If the value of *file* is 0, *dlopen()* shall provide a *handle* on a global symbol object. This object shall  
 8513 provide access to the symbols from an ordered set of objects consisting of the original program  
 8514 image file, together with any objects loaded at program start-up as specified by that process  
 8515 image file (for example, shared libraries), and the set of objects loaded using a *dlopen()* operation  
 8516 together with the RTLD\_GLOBAL flag. As the latter set of objects can change during execution,  
 8517 the set identified by *handle* can also change dynamically.

8518 Only a single copy of an object file is brought into the address space, even if *dlopen()* is invoked  
 8519 multiple times in reference to the file, and even if different pathnames are used to reference the  
 8520 file.

8521 The *mode* parameter describes how *dlopen()* shall operate upon *file* with respect to the processing  
 8522 of relocations and the scope of visibility of the symbols provided within *file*. When an object is  
 8523 brought into the address space of a process, it may contain references to symbols whose  
 8524 addresses are not known until the object is loaded. These references shall be relocated before the  
 8525 symbols can be accessed. The *mode* parameter governs when these relocations take place and  
 8526 may have the following values:

8527 RTLD\_LAZY Relocations shall be performed at an implementation-defined time,  
 8528 ranging from the time of the *dlopen()* call until the first reference to a  
 8529 given symbol occurs. Specifying RTLD\_LAZY should improve  
 8530 performance on implementations supporting dynamic symbol binding as  
 8531 a process may not reference all of the functions in any given object. And,  
 8532 for systems supporting dynamic symbol resolution for normal process  
 8533 execution, this behavior mimics the normal handling of process  
 8534 execution.

8535 RTLD\_NOW All necessary relocations shall be performed when the object is first  
 8536 loaded. This may waste some processing if relocations are performed for  
 8537 functions that are never referenced. This behavior may be useful for  
 8538 applications that need to know as soon as an object is loaded that all

8539 symbols referenced during execution are available.

8540 Any object loaded by *dlopen()* that requires relocations against global symbols can reference the  
8541 symbols in the original process image file, any objects loaded at program start-up, from the  
8542 object itself as well as any other object included in the same *dlopen()* invocation, and any objects  
8543 that were loaded in any *dlopen()* invocation and which specified the RTLD\_GLOBAL flag. To  
8544 determine the scope of visibility for the symbols loaded with a *dlopen()* invocation, the *mode*  
8545 parameter should be a bitwise-inclusive OR with one of the following values:

8546 RTLD\_GLOBAL The object's symbols shall be made available for the relocation processing  
8547 of any other object. In addition, symbol lookup using *dlopen(0, mode)* and  
8548 an associated *dlsym()* allows objects loaded with this *mode* to be searched.

8549 RTLD\_LOCAL The object's symbols shall not be made available for the relocation  
8550 processing of any other object.

8551 If neither RTLD\_GLOBAL nor RTLD\_LOCAL are specified, then an implementation-defined  
8552 default behavior shall be applied.

8553 If a *file* is specified in multiple *dlopen()* invocations, *mode* is interpreted at each invocation. Note,  
8554 however, that once RTLD\_NOW has been specified all relocations shall have been completed  
8555 rendering further RTLD\_NOW operations redundant and any further RTLD\_LAZY operations  
8556 irrelevant. Similarly, note that once RTLD\_GLOBAL has been specified the object shall maintain  
8557 the RTLD\_GLOBAL status regardless of any previous or future specification of RTLD\_LOCAL,  
8558 as long as the object remains in the address space (see *dlclose()*).

8559 Symbols introduced into a program through calls to *dlopen()* may be used in relocation  
8560 activities. Symbols so introduced may duplicate symbols already defined by the program or  
8561 previous *dlopen()* operations. To resolve the ambiguities such a situation might present, the  
8562 resolution of a symbol reference to symbol definition is based on a symbol resolution order. Two  
8563 such resolution orders are defined: *load* or *dependency* ordering. Load order establishes an  
8564 ordering among symbol definitions, such that the definition first loaded (including definitions  
8565 from the image file and any dependent objects loaded with it) has priority over objects added  
8566 later (via *dlopen()*). Load ordering is used in relocation processing. Dependency ordering uses a  
8567 breadth-first order starting with a given object, then all of its dependencies, then any dependents  
8568 of those, iterating until all dependencies are satisfied. With the exception of the global symbol  
8569 object obtained via a *dlopen()* operation on a *file* of 0, dependency ordering is used by the  
8570 *dlsym()* function. Load ordering is used in *dlsym()* operations upon the global symbol object.

8571 When an object is first made accessible via *dlopen()* it and its dependent objects are added in  
8572 dependency order. Once all the objects are added, relocations are performed using load order.  
8573 Note that if an object or its dependencies had been previously loaded, the load and dependency  
8574 orders may yield different resolutions.

8575 The symbols introduced by *dlopen()* operations, and available through *dlsym()* are at a  
8576 minimum those which are exported as symbols of global scope by the object. Typically such  
8577 symbols shall be those that were specified in (for example) C source code as having *extern*  
8578 linkage. The precise manner in which an implementation constructs the set of exported symbols  
8579 for a *dlopen()* object is specified by that implementation.

8580 **RETURN VALUE**

8581 If *file* cannot be found, cannot be opened for reading, is not of an appropriate object format for  
8582 processing by *dlopen()*, or if an error occurs during the process of loading *file* or relocating its  
8583 symbolic references, *dlopen()* shall return NULL. More detailed diagnostic information shall be  
8584 available through *dlerror()*.

8585 **ERRORS**

8586       No errors are defined.

8587 **EXAMPLES**

8588       None.

8589 **APPLICATION USAGE**

8590       None.

8591 **RATIONALE**

8592       None.

8593 **FUTURE DIRECTIONS**

8594       None.

8595 **SEE ALSO**

8596       *dlclose()*, *dLError()*, *dlsym()*, the Base Definitions volume of IEEE Std 1003.1-200x, <**dlfcn.h**>

8597 **CHANGE HISTORY**

8598       First released in Issue 5.

8599 **NAME**8600 `dlsym` — obtain the address of a symbol from a `dlopen()` object8601 **SYNOPSIS**8602 XSI 

```
#include <dlfcn.h>
```

8603 

```
void *dlsym(void *restrict handle, const char *restrict name);
```

8604

8605 **DESCRIPTION**

8606 The `dlsym()` function shall obtain the address of a symbol defined within an object made |  
 8607 accessible through a `dlopen()` call. The `handle` argument is the value returned from a call to |  
 8608 `dlopen()` (and which has not since been released via a call to `dlclose()`), and `name` is the symbol's |  
 8609 name as a character string.

8610 The `dlsym()` function shall search for the named symbol in all objects loaded automatically as a |  
 8611 result of loading the object referenced by `handle` (see `dlopen()`). Load ordering is used in `dlsym()` |  
 8612 operations upon the global symbol object. The symbol resolution algorithm used shall be |  
 8613 dependency order as described in `dlopen()`.

8614 The `RTLD_NEXT` flag is reserved for future use.8615 **RETURN VALUE**

8616 If `handle` does not refer to a valid object opened by `dlopen()`, or if the named symbol cannot be |  
 8617 found within any of the objects associated with `handle`, `dlsym()` shall return NULL. More |  
 8618 detailed diagnostic information shall be available through `dlerror()`.

8619 **ERRORS**

8620 No errors are defined.

8621 **EXAMPLES**

8622 The following example shows how `dlopen()` and `dlsym()` can be used to access either function or |  
 8623 data objects. For simplicity, error checking has been omitted.

```
8624 void *handle;
8625 int *iptr, (*fptr)(int);

8626 /* open the needed object */
8627 handle = dlopen("/usr/home/me/libfoo.so", RTLD_LOCAL | RTLD_LAZY);

8628 /* find the address of function and data objects */
8629 fptr = (int (*)(int))dlsym(handle, "my_function");
8630 iptr = (int *)dlsym(handle, "my_object");

8631 /* invoke function, passing value of integer as a parameter */
8632 (*fptr)(*iptr);
```

8633 **APPLICATION USAGE**

8634 Special purpose values for `handle` are reserved for future use. These values and their meanings |  
 8635 are:

8636 **RTLD\_DEFAULT** The symbol lookup happens in the normal global scope; that is, a search for a |  
 8637 symbol using this handle would find the same definition as a direct use of this |  
 8638 symbol in the program code.

8639 **RTLD\_NEXT** Specifies the next object after this one that defines `name`. *This one* refers to the |  
 8640 object containing the invocation of `dlsym()`. The *next* object is the one found |  
 8641 upon the application of a load order symbol resolution algorithm (see |  
 8642 `dlopen()`). The next object is either one of global scope (because it was |  
 8643 introduced as part of the original process image or because it was added with

8644 a *dlopen()* operation including the `RTLD_GLOBAL` flag), or is an object that  
8645 was included in the same *dlopen()* operation that loaded this one.

8646 The `RTLD_NEXT` flag is useful to navigate an intentionally created hierarchy  
8647 of multiply-defined symbols created through *interposition*. For example, if a  
8648 program wished to create an implementation of *malloc()* that embedded some  
8649 statistics gathering about memory allocations, such an implementation could  
8650 use the real *malloc()* definition to perform the memory allocation—and itself  
8651 only embed the necessary logic to implement the statistics gathering function.

8652 **RATIONALE**  
8653 None.

8654 **FUTURE DIRECTIONS**  
8655 None.

8656 **SEE ALSO**  
8657 *dlclose()*, *dLError()*, *dlopen()*, the Base Definitions volume of IEEE Std 1003.1-200x, <**dlfcn.h**>

8658 **CHANGE HISTORY**  
8659 First released in Issue 5.

8660 **Issue 6**  
8661 The **restrict** keyword is added to the *dlsym()* prototype for alignment with the  
8662 ISO/IEC 9899:1999 standard.

## 8663 NAME

8664 drand48, erand48, jrand48, lcong48, lrand48, mrand48, nrand48, seed48, srand48 — generate  
8665 uniformly distributed pseudo-random numbers

## 8666 SYNOPSIS

```
8667 xSI #include <stdlib.h>
8668
8668 double drand48(void);
8669 double erand48(unsigned short xsubi[3]);
8670 long jrand48(unsigned short xsubi[3]);
8671 void lcong48(unsigned short param[7]);
8672 long lrand48(void);
8673 long mrand48(void);
8674 long nrand48(unsigned short xsubi[3]);
8675 unsigned short *seed48(unsigned short seed16v[3]);
8676 void srand48(long seedval);
8677
```

## 8678 DESCRIPTION

8679 This family of functions shall generate pseudo-random numbers using a linear congruential  
8680 algorithm and 48-bit integer arithmetic.

8681 The *drand48()* and *erand48()* functions shall return non-negative, double-precision, floating-  
8682 point values, uniformly distributed over the interval [0.0,1.0).

8683 The *lrnd48()* and *nrnd48()* functions shall return non-negative, long integers, uniformly  
8684 distributed over the interval  $[0, 2^{31})$ .

8685 The *mrnd48()* and *jrnd48()* functions shall return signed long integers uniformly distributed  
8686 over the interval  $[-2^{31}, 2^{31})$ .

8687 The *srand48()*, *seed48()*, and *lcong48()* are initialization entry points, one of which should be  
8688 invoked before either *drand48()*, *lrnd48()*, or *mrnd48()* is called. (Although it is not  
8689 recommended practice, constant default initializer values shall be supplied automatically if  
8690 *drand48()*, *lrnd48()*, or *mrnd48()* is called without a prior call to an initialization entry point.)  
8691 The *erand48()*, *nrnd48()*, and *jrnd48()* functions do not require an initialization entry point to  
8692 be called first.

8693 All the routines work by generating a sequence of 48-bit integer values,  $X_i$ , according to the  
8694 linear congruential formula:

$$8695 \quad X_{n+1} = (aX_n + c)_{\text{mod } m} \quad n \geq 0$$

8696 The parameter  $m = 2^{48}$ ; hence 48-bit integer arithmetic is performed. Unless *lcong48()* is invoked,  
8697 the multiplier value  $a$  and the addend value  $c$  are given by:

$$8698 \quad a = 5\text{DEECE66D}_{16} = 273673163155_8$$

$$8699 \quad c = \text{B}_{16} = 13_8$$

8700 The value returned by any of the *drand48()*, *erand48()*, *jrnd48()*, *lrnd48()*, *mrnd48()*, or  
8701 *nrnd48()* functions is computed by first generating the next 48-bit  $X_i$  in the sequence. Then the  
8702 appropriate number of bits, according to the type of data item to be returned, are copied from  
8703 the high-order (leftmost) bits of  $X_i$  and transformed into the returned value.

8704 The *drand48()*, *lrnd48()*, and *mrnd48()* functions store the last 48-bit  $X_i$  generated in an  
8705 internal buffer; that is why the application shall ensure that these are initialized prior to being  
8706 invoked. The *erand48()*, *nrnd48()*, and *jrnd48()* functions require the calling program to  
8707 provide storage for the successive  $X_i$  values in the array specified as an argument when the

8708 functions are invoked. That is why these routines do not have to be initialized; the calling  
8709 program merely has to place the desired initial value of  $X_i$  into the array and pass it as an  
8710 argument. By using different arguments, *erand48()*, *rand48()*, and *jrand48()* allow separate  
8711 modules of a large program to generate several *independent* streams of pseudo-random numbers;  
8712 that is, the sequence of numbers in each stream shall *not* depend upon how many times the  
8713 routines are called to generate numbers for the other streams.

8714 The initializer function *srand48()* sets the high-order 32 bits of  $X_i$  to the low-order 32 bits  
8715 contained in its argument. The low-order 16 bits of  $X_i$  are set to the arbitrary value  $330E_{16}$ .

8716 The initializer function *seed48()* sets the value of  $X_i$  to the 48-bit value specified in the argument  
8717 array. The low-order 16 bits of  $X_i$  are set to the low-order 16 bits of *seed16v[0]*. The mid-order 16  
8718 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  
8719 low-order 16 bits of *seed16v[2]*. In addition, the previous value of  $X_i$  is copied into a 48-bit  
8720 internal buffer, used only by *seed48()*, and a pointer to this buffer is the value returned by  
8721 *seed48()*. This returned pointer, which can just be ignored if not needed, is useful if a program is  
8722 to be restarted from a given point at some future time—use the pointer to get at and store the  
8723 last  $X_i$  value, and then use this value to reinitialize via *seed48()* when the program is restarted.

8724 The initializer function *lcg48()* allows the user to specify the initial  $X_i$ , the multiplier value *a*,  
8725 and the addend value *c*. Argument array elements *param[0-2]* specify  $X_i$ , *param[3-5]* specify the  
8726 multiplier *a*, and *param[6]* specifies the 16-bit addend *c*. After *lcg48()* is called, a subsequent  
8727 call to either *srand48()* or *seed48()* shall restore the standard multiplier and addend values, *a* and  
8728 *c*, specified above.

8729 The *drand48()*, *lrnd48()*, and *mrnd48()* functions need not be reentrant. A function that is not  
8730 required to be reentrant is not required to be thread-safe.

#### 8731 RETURN VALUE

8732 As described in the DESCRIPTION above.

#### 8733 ERRORS

8734 No errors are defined.

#### 8735 EXAMPLES

8736 None.

#### 8737 APPLICATION USAGE

8738 None.

#### 8739 RATIONALE

8740 None.

#### 8741 FUTURE DIRECTIONS

8742 None.

#### 8743 SEE ALSO

8744 *rand()*, the Base Definitions volume of IEEE Std 1003.1-200x, <stdlib.h>

#### 8745 CHANGE HISTORY

8746 First released in Issue 1. Derived from Issue 1 of the SVID.

#### 8747 Issue 5

8748 A note indicating that these functions need not be reentrant is added to the DESCRIPTION.

#### 8749 Issue 6

8750 The DESCRIPTION is updated to avoid use of the term “must” for application requirements.



8751 **NAME**

8752 dup, dup2 — duplicate an open file descriptor

8753 **SYNOPSIS**

8754 #include &lt;unistd.h&gt;

8755 int dup(int *fil-des*);8756 int dup2(int *fil-des*, int *fil-des2*);8757 **DESCRIPTION**8758 The *dup()* and *dup2()* functions provide an alternative interface to the service provided by *fcntl()* using the `F_DUPFD` command. The call:8760 *fid* = dup(*fil-des*);

8761 shall be equivalent to:

8762 *fid* = fcntl(*fil-des*, `F_DUPFD`, 0);

8763 The call:

8764 *fid* = dup2(*fil-des*, *fil-des2*);

8765 shall be equivalent to:

8766 close(*fil-des2*);8767 *fid* = fcntl(*fil-des*, `F_DUPFD`, *fil-des2*);

8768 except for the following:

- 8769 • If *fil-des2* is less than 0 or greater than or equal to `{OPEN_MAX}`, *dup2()* shall return `-1` with *errno* set to `[EBADF]`.
- 8770
- 8771 • If *fil-des* is a valid file descriptor and is equal to *fil-des2*, *dup2()* shall return *fil-des2* without closing it.
- 8772
- 8773 • If *fil-des* is not a valid file descriptor, *dup2()* shall return `-1` and shall not close *fil-des2*.
- 8774 • The value returned shall be equal to the value of *fil-des2* upon successful completion, or `-1` upon failure.
- 8775

8776 **RETURN VALUE**8777 Upon successful completion a non-negative integer, namely the file descriptor, shall be returned; otherwise, `-1` shall be returned and *errno* set to indicate the error.8779 **ERRORS**8780 The *dup()* function shall fail if:8781 `[EBADF]` The *fil-des* argument is not a valid open file descriptor.8782 `[EMFILE]` The number of file descriptors in use by this process would exceed `{OPEN_MAX}`.8784 The *dup2()* function shall fail if:8785 `[EBADF]` The *fil-des* argument is not a valid open file descriptor or the argument *fil-des2* is negative or greater than or equal to `{OPEN_MAX}`.8787 `[EINTR]` The *dup2()* function was interrupted by a signal.

8788 **EXAMPLES**8789 **Redirecting Standard Output to a File**

8790 The following example closes standard output for the current processes, re-assigns standard  
8791 output to go to the file referenced by *pfid*, and closes the original file descriptor to clean up.

```
8792 #include <unistd.h>  
8793 ...  
8794 int pfd;  
8795 ...  
8796 close(1);  
8797 dup(pfd);  
8798 close(pfd);  
8799 ...
```

8800 **Redirecting Error Messages**

8801 The following example redirects messages from *stderr* to *stdout*.

```
8802 #include <unistd.h>  
8803 ...  
8804 dup2(1, 2);  
8805 ...
```

8806 **APPLICATION USAGE**

8807 None.

8808 **RATIONALE**

8809 The *dup()* and *dup2()* functions are redundant. Their services are also provided by the *fcntl()*  
8810 function. They have been included in this volume of IEEE Std 1003.1-200x primarily for historical  
8811 reasons, since many existing applications use them.

8812 While the brief code segment shown is very similar in behavior to *dup2()*, a conforming  
8813 implementation based on other functions defined in this volume of IEEE Std 1003.1-200x  
8814 is significantly more complex. Least obvious is the possible effect of a signal-catching function that  
8815 could be invoked between steps and allocate or deallocate file descriptors. This could be avoided  
8816 by blocking signals.

8817 The *dup2()* function is not marked obsolescent because it presents a type-safe version of  
8818 functionality provided in a type-unsafe version by *fcntl()*. It is used in the POSIX Ada binding.

8819 The *dup2()* function is not intended for use in critical regions as a synchronization mechanism.

8820 In the description of [EBADF], the case of *fildes* being out of range is covered by the given case of  
8821 *fildes* not being valid. The descriptions for *fildes* and *fildes2* are different because the only kind of  
8822 invalidity that is relevant for *fildes2* is whether it is out of range; that is, it does not matter  
8823 whether *fildes2* refers to an open file when the *dup2()* call is made.

8824 **FUTURE DIRECTIONS**

8825 None.

8826 **SEE ALSO**

8827 *close()*, *fcntl()*, *open()*, the Base Definitions volume of IEEE Std 1003.1-200x, <**unistd.h**>

8828 **CHANGE HISTORY**

8829 First released in Issue 1. Derived from Issue 1 of the SVID.

8830 **Issue 6**

8831 The RATIONALE section is added.

## 8832 NAME

8833 ecvt, fcvt, gcvt — convert a floating-point number to a string (**LEGACY**)

## 8834 SYNOPSIS

```
8835 xSI #include <stdlib.h>
8836
8836 char *ecvt(double value, int ndigit, int *restrict decpt,
8837           int *restrict sign);
8838 char *fcvt(double value, int ndigit, int *restrict decpt,
8839           int *restrict sign);
8840 char *gcvt(double value, int ndigit, char *buf);
8841
```

## 8842 DESCRIPTION

8843 The *ecvt()*, *fcvt()*, and *gcvt()* functions shall convert floating-point numbers to null-terminated strings.  
8844

8845 The *ecvt()* function shall convert *value* to a null-terminated string of *ndigit* digits (where *ndigit* is  
8846 reduced to an unspecified limit determined by the precision of a **double**) and return a pointer to  
8847 the string. The high-order digit shall be non-zero, unless the value is 0. The low-order digit shall  
8848 be rounded in an implementation-defined manner. The position of the radix character relative to  
8849 the beginning of the string shall be stored in the integer pointed to by *decpt* (negative means to  
8850 the left of the returned digits). If *value* is zero, it is unspecified whether the integer pointed to by  
8851 *decpt* would be 0 or 1. The radix character shall not be included in the returned string. If the sign  
8852 of the result is negative, the integer pointed to by *sign* shall be non-zero; otherwise, it shall be 0.

8853 If the converted value is out of range or is not representable, the contents of the returned string  
8854 are unspecified.

8855 The *fcvt()* function shall be equivalent to *ecvt()*, except that *ndigit* specifies the number of digits  
8856 desired after the radix character. The total number of digits in the result string is restricted to an  
8857 unspecified limit as determined by the precision of a **double**.

8858 The *gcvt()* function shall convert *value* to a null-terminated string (similar to that of the %g  
8859 conversion specification format of *printf()*) in the array pointed to by *buf* and shall return *buf*. It  
8860 shall produce *ndigit* significant digits (limited to an unspecified value determined by the  
8861 precision of a **double**) in the %E conversion specification format of *printf()* if possible, or the %e  
8862 conversion specification format of *printf()* (scientific notation) otherwise. A minus sign shall be  
8863 included in the returned string if *value* is less than 0. A radix character shall be included in the  
8864 returned string if *value* is not a whole number. Trailing zeros shall be suppressed where *value* is  
8865 not a whole number. The radix character is determined by the current locale. If *setlocale()* has not  
8866 been called successfully, the default locale, POSIX, is used. The default locale specifies a period  
8867 ('.') as the radix character. The *LC\_NUMERIC* category determines the value of the radix  
8868 character within the current locale.

8869 These functions need not be reentrant. A function that is not required to be reentrant is not  
8870 required to be thread-safe.

## 8871 RETURN VALUE

8872 The *ecvt()* and *fcvt()* functions shall return a pointer to a null-terminated string of digits.8873 The *gcvt()* function shall return *buf*.8874 The return values from *ecvt()* and *fcvt()* may point to static data which may be overwritten by  
8875 subsequent calls to these functions.

8876 **ERRORS**

8877 No errors are defined.

8878 **EXAMPLES**

8879 None.

8880 **APPLICATION USAGE**8881 *sprintf()* is preferred over this function.8882 **RATIONALE**

8883 None.

8884 **FUTURE DIRECTIONS**

8885 These functions may be withdrawn in a future version.

8886 **SEE ALSO**8887 *printf()*, *setlocale()*, the Base Definitions volume of IEEE Std 1003.1-200x, <stdlib.h>8888 **CHANGE HISTORY**

8889 First released in Issue 4, Version 2.

8890 **Issue 5**

8891 Moved from X/OPEN UNIX extension to BASE.

8892 Normative text previously in the APPLICATION USAGE section is moved to the  
8893 DESCRIPTION.

8894 A note indicating that these functions need not be reentrant is added to the DESCRIPTION.

8895 **Issue 6**

8896 In the DESCRIPTION, the note about reentrancy is expanded to cover thread-safety.

8897 This function is marked LEGACY.

8898 The **restrict** keyword is added to the *ecvt()* and *fcvt()* prototypes for alignment with the  
8899 ISO/IEC 9899:1999 standard.8900 The DESCRIPTION is updated to explicitly use “conversion specification” to describe %g, %f,  
8901 and %e.

8902 **NAME**8903 encrypt — encoding function (**CRYPT**)8904 **SYNOPSIS**8905 XSI `#include <unistd.h>`8906 `void encrypt(char block[64], int edflag);`

8907

8908 **DESCRIPTION**8909 The *encrypt()* function shall provide access to an implementation-defined encoding algorithm. |8910 The key generated by *setkey()* is used to encrypt the string *block* with *encrypt()*. |8911 The *block* argument to *encrypt()* shall be an array of length 64 bytes containing only the bytes |

8912 with values of 0 and 1. The array is modified in place to a similar array using the key set by |

8913 *setkey()*. If *edflag* is 0, the argument is encoded. If *edflag* is 1, the argument may be decoded (see |8914 the APPLICATION USAGE section); if the argument is not decoded, *errno* shall be set to |

8915 [ENOSYS].

8916 The *encrypt()* function shall not change the setting of *errno* if successful. An application wishing |8917 to check for error situations should set *errno* to 0 before calling *encrypt()*. If *errno* is non-zero on |

8918 return, an error has occurred.

8919 The *encrypt()* function need not be reentrant. A function that is not required to be reentrant is |

8920 not required to be thread-safe.

8921 **RETURN VALUE**8922 The *encrypt()* function shall not return a value.8923 **ERRORS**8924 The *encrypt()* function shall fail if:

8925 [ENOSYS] The functionality is not supported on this implementation.

8926 **EXAMPLES**

8927 None.

8928 **APPLICATION USAGE**8929 Historical implementations of the *encrypt()* function used a rather primitive encoding algorithm. |

8930 In some environments, decoding might not be implemented. This is related to some Government |

8931 restrictions on encryption and decryption routines. Historical practice has been to ship a |

8932 different version of the encryption library without the decryption feature in the routines |

8933 supplied. Thus the exported version of *encrypt()* does encoding but not decoding.8934 **RATIONALE**

8935 None.

8936 **FUTURE DIRECTIONS**

8937 None.

8938 **SEE ALSO**8939 *crypt()*, *setkey()*, the Base Definitions volume of IEEE Std 1003.1-200x, `<unistd.h>`8940 **CHANGE HISTORY**

8941 First released in Issue 1. Derived from Issue 1 of the SVID.

8942 **Issue 5**

8943 A note indicating that this function need not be reentrant is added to the DESCRIPTION.

8944 **Issue 6**

8945 In the DESCRIPTION, the note about reentrancy is expanded to cover thread-safety.

8946 **NAME**

8947 endgrent, getgrent, setgrent — group database entry functions

8948 **SYNOPSIS**

```
8949 xsi #include <grp.h>
8950 void endgrent(void);
8951 struct group *getgrent(void);
8952 void setgrent(void);
8953
```

8954 **DESCRIPTION**

8955 The *getgrent()* function shall return a pointer to a structure containing the broken-out fields of an  
8956 entry in the group database. When first called, *getgrent()* shall return a pointer to a **group**  
8957 structure containing the first entry in the group database. Thereafter, it shall return a pointer to a  
8958 **group** structure containing the next group structure in the group database, so successive calls  
8959 may be used to search the entire database.

8960 An implementation that provides extended security controls may impose further  
8961 implementation-defined restrictions on accessing the group database. In particular, the system  
8962 may deny the existence of some or all of the group database entries associated with groups other  
8963 than those groups associated with the caller and may omit users other than the caller from the  
8964 list of members of groups in database entries that are returned.

8965 The *setgrent()* function shall rewind the group database to allow repeated searches.

8966 The *endgrent()* function may be called to close the group database when processing is complete.

8967 These functions need not be reentrant. A function that is not required to be reentrant is not  
8968 required to be thread-safe.

8969 **RETURN VALUE**

8970 When first called, *getgrent()* shall return a pointer to the first group structure in the group  
8971 database. Upon subsequent calls it shall return the next group structure in the group database.  
8972 The *getgrent()* function shall return a null pointer on end-of-file or an error and *errno* may be set  
8973 to indicate the error.

8974 The return value may point to a static area which is overwritten by a subsequent call to  
8975 *getgrgid()*, *getgrnam()*, or *getgrent()*.

8976 **ERRORS**

8977 The *getgrent()* function may fail if:

- |      |          |  |
|------|----------|--|
| 8978 | [EINTR]  | A signal was caught during the operation.                              |
| 8979 | [EIO]    | An I/O error has occurred.   |
| 8980 | [EMFILE] | {OPEN_MAX} file descriptors are currently open in the calling process. |
| 8981 | [ENFILE] | The maximum allowable number of files is currently open in the system. |



**8982 EXAMPLES**

8983 None.

**8984 APPLICATION USAGE**

8985 These functions are provided due to their historical usage. Applications should avoid  
8986 dependencies on fields in the group database, whether the database is a single file, or where in  
8987 the file system name space the database resides. Applications should use *getgrnam()* and  
8988 *getgrgid()* whenever possible because it avoids these dependencies.

**8989 RATIONALE**

8990 None.

**8991 FUTURE DIRECTIONS**

8992 None.

**8993 SEE ALSO**

8994 *getgrgid()*, *getgrnam()*, *getlogin()*, *getpwent()*, the Base Definitions volume of  
8995 IEEE Std 1003.1-200x, <grp.h>

**8996 CHANGE HISTORY**

8997 First released in Issue 4, Version 2.

**8998 Issue 5**

8999 Moved from X/OPEN UNIX extension to BASE.

9000 Normative text previously in the APPLICATION USAGE section is moved to the RETURN  
9001 VALUE section.

9002 A note indicating that these functions need not be reentrant is added to the DESCRIPTION.

**9003 Issue 6**

9004 In the DESCRIPTION, the note about reentrancy is expanded to cover thread-safety.

9005 **NAME**

9006 endhostent, gethostent, sethostent — network host database functions

9007 **SYNOPSIS**

9008 #include &lt;netdb.h&gt;

9009 void endhostent(void);

9010 struct hostent \*gethostent(void);

9011 void sethostent(int *stayopen*);9012 **DESCRIPTION**

9013 These functions shall retrieve information about hosts. This information is considered to be |  
9014 stored in a database that can be accessed sequentially or randomly. The implementation of this |  
9015 database is unspecified. |

9016 **Note:** In many cases it is implemented by the Domain Name System, as documented in RFC 1034,  
9017 RFC 1035, and RFC 1886.

9018 The *sethostent()* function shall open a connection to the database and set the next entry for  
9019 retrieval to the first entry in the database. If the *stayopen* argument is non-zero, the connection  
9020 shall not be closed by a call to *gethostent()*, *gethostbyname()*, or *gethostbyaddr()*, and the  
9021 implementation may maintain an open file descriptor.

9022 The *gethostent()* function shall read the next entry in the database, opening and closing a  
9023 connection to the database as necessary.

9024 Entries shall be returned in **hostent** structures. Refer to *gethostbyaddr()* for a definition of the |  
9025 **hostent** structure.

9026 The *endhostent()* function shall close the connection to the database, releasing any open file  
9027 descriptor.

9028 These functions need not be reentrant. A function that is not required to be reentrant is not  
9029 required to be thread-safe.

9030 **RETURN VALUE**

9031 Upon successful completion, the *gethostent()* function shall return a pointer to a **hostent**  
9032 structure if the requested entry was found, and a null pointer if the end of the database was  
9033 reached or the requested entry was not found.

9034 **ERRORS**9035 No errors are defined for *endhostent()*, *gethostent()*, and *sethostent()*.9036 **EXAMPLES**

9037 None.

9038 **APPLICATION USAGE**

9039 The *gethostent()* function may return pointers to static data, which may be overwritten by  
9040 subsequent calls to any of these functions.

9041 **RATIONALE**

9042 None.

9043 **FUTURE DIRECTIONS**

9044 None.

9045 **SEE ALSO**

9046 *endservent()*, *gethostbyaddr()*, *gethostbyname()*, the Base Definitions volume of  
9047 IEEE Std 1003.1-200x, <netdb.h>

9048 **CHANGE HISTORY**

9049 First released in Issue 6. Derived from the XNS, Issue 5.2 specification.

9050 **NAME**

9051 endnetent, getnetbyaddr, getnetbyname, getnetent, setnetent — network database functions

9052 **SYNOPSIS**

9053 #include &lt;netdb.h&gt;

9054 void endnetent(void);

9055 struct netent \*getnetbyaddr(uint32\_t net, int type);

9056 struct netent \*getnetbyname(const char \*name);

9057 struct netent \*getnetent(void);

9058 void setnetent(int stayopen);

9059 **DESCRIPTION**

9060 These functions shall retrieve information about networks. This information is considered to be |  
9061 stored in a database that can be accessed sequentially or randomly. The implementation of this |  
9062 database is unspecified. |

9063 The *setnetent()* function shall open and rewind the database. If the *stayopen* argument is non-  
9064 zero, the connection to the *net* database shall not be closed after each call to *getnetent()* (either  
9065 directly, or indirectly through one of the other *getnet\**(*)* functions), and the implementation may  
9066 maintain an open file descriptor to the database.

9067 The *getnetent()* function shall read the next entry of the database, opening and closing a  
9068 connection to the database as necessary.

9069 The *getnetbyaddr()* function shall search the database from the beginning, and find the first entry  
9070 for which the address family specified by *type* matches the *n\_addrtype* member and the network  
9071 number *net* matches the *n\_net* member, opening and closing a connection to the database as  
9072 necessary. The *net* argument shall be the network number in host byte order.

9073 The *getnetbyname()* function shall search the database from the beginning and find the first entry  
9074 for which the network name specified by *name* matches the *n\_name* member, opening and  
9075 closing a connection to the database as necessary.

9076 The *getnetbyaddr()*, *getnetbyname()*, and *getnetent()*, functions shall each return a pointer to a  
9077 **netent** structure, the members of which shall contain the fields of an entry in the network  
9078 database.

9079 The *endnetent()* function shall close the database, releasing any open file descriptor.

9080 These functions need not be reentrant. A function that is not required to be reentrant is not  
9081 required to be thread-safe.

9082 **RETURN VALUE**

9083 Upon successful completion, *getnetbyaddr()*, *getnetbyname()*, and *getnetent()*, shall return a  
9084 pointer to a **netent** structure if the requested entry was found, and a null pointer if the end of the  
9085 database was reached or the requested entry was not found. Otherwise, a null pointer shall be  
9086 returned.

9087 **ERRORS**

9088 No errors are defined.

9089 **EXAMPLES**

9090       None.

9091 **APPLICATION USAGE**9092       The *getnetbyaddr()*, *getnetbyname()*, and *getnetent()*, functions may return pointers to static data,  
9093       which may be overwritten by subsequent calls to any of these functions.9094 **RATIONALE**

9095       None.

9096 **FUTURE DIRECTIONS**

9097       None.

9098 **SEE ALSO**9099       The Base Definitions volume of IEEE Std 1003.1-200x, <**netdb.h**>9100 **CHANGE HISTORY**

9101       First released in Issue 6. Derived from the XNS, Issue 5.2 specification.

9102 **NAME**

9103 endprotoent, getprotobyname, getprotobynumber, getprotoent, setprotoent — network protocol  
9104 database functions

9105 **SYNOPSIS**

```
9106 #include <netdb.h>
9107
9108 void endprotoent(void);
9109 struct protoent *getprotobyname(const char *name);
9110 struct protoent *getprotobynumber(int proto);
9111 struct protoent *getprotoent(void);
9112 void setprotoent(int stayopen);
```

9112 **DESCRIPTION**

9113 These functions shall retrieve information about protocols. This information is considered to be |  
9114 stored in a database that can be accessed sequentially or randomly. The implementation of this |  
9115 database is unspecified. |

9116 The *setprotoent()* function shall open a connection to the database, and set the next entry to the  
9117 first entry. If the *stayopen* argument is non-zero, the connection to the network protocol database  
9118 shall not be closed after each call to *getprotoent()* (either directly, or indirectly through one of the  
9119 other *getproto\**() functions), and the implementation may maintain an open file descriptor for  
9120 the database.

9121 The *getprotobyname()* function shall search the database from the beginning and find the first  
9122 entry for which the protocol name specified by *name* matches the *p\_name* member, opening and  
9123 closing a connection to the database as necessary.

9124 The *getprotobynumber()* function shall search the database from the beginning and find the first  
9125 entry for which the protocol number specified by *proto* matches the *p\_proto* member, opening  
9126 and closing a connection to the database as necessary.

9127 The *getprotoent()* function shall read the next entry of the database, opening and closing a  
9128 connection to the database as necessary.

9129 The *getprotobyname()*, *getprotobynumber()*, and *getprotoent()*, functions shall each return a pointer  
9130 to a **protoent** structure, the members of which shall contain the fields of an entry in the network  
9131 protocol database.

9132 The *endprotoent()* function shall close the connection to the database, releasing any open file  
9133 descriptor.

9134 These functions need not be reentrant. A function that is not required to be reentrant is not  
9135 required to be thread-safe.

9136 **RETURN VALUE**

9137 Upon successful completion, *getprotobyname()*, *getprotobynumber()*, and *getprotoent()* return a  
9138 pointer to a **protoent** structure if the requested entry was found, and a null pointer if the end of  
9139 the database was reached or the requested entry was not found. Otherwise, a null pointer is  
9140 returned.

9141 **ERRORS**

9142 No errors are defined.

9143 **EXAMPLES**

9144       None.

9145 **APPLICATION USAGE**9146       The *getprotobyname()*, *getprotobynumber()*, and *getprotoent()* functions may return pointers to  
9147       static data, which may be overwritten by subsequent calls to any of these functions.9148 **RATIONALE**

9149       None.

9150 **FUTURE DIRECTIONS**

9151       None.

9152 **SEE ALSO**9153       The Base Definitions volume of IEEE Std 1003.1-200x, <**netdb.h**>9154 **CHANGE HISTORY**

9155       First released in Issue 6. Derived from the XNS, Issue 5.2 specification.

9156 **NAME**

9157 endpwent, getpwent, setpwent — user database functions

9158 **SYNOPSIS**

```
9159 xSI #include <pwd.h>
9160 void endpwent(void);
9161 struct passwd *getpwent(void);
9162 void setpwent(void);
9163
```

9164 **DESCRIPTION**

9165 These functions shall retrieve information about users. |

9166 The *getpwent()* function shall return a pointer to a structure containing the broken-out fields of |  
 9167 an entry in the user database. Each entry in the user database contains a **passwd** structure. When |  
 9168 first called, *getpwent()* shall return a pointer to a **passwd** structure containing the first entry in |  
 9169 the user database. Thereafter, it shall return a pointer to a **passwd** structure containing the next |  
 9170 entry in the user database. Successive calls can be used to search the entire user database.

9171 If an end-of-file or an error is encountered on reading, *getpwent()* shall return a null pointer.

9172 An implementation that provides extended security controls may impose further |  
 9173 implementation-defined restrictions on accessing the user database. In particular, the system |  
 9174 may deny the existence of some or all of the user database entries associated with users other |  
 9175 than the caller.

9176 The *setpwent()* function effectively rewinds the user database to allow repeated searches.9177 The *endpwent()* function may be called to close the user database when processing is complete.

9178 These functions need not be reentrant. A function that is not required to be reentrant is not |  
 9179 required to be thread-safe.

9180 **RETURN VALUE**9181 The *getpwent()* function shall return a null pointer on end-of-file or error.9182 **ERRORS**9183 The *getpwent()*, *setpwent()*, and *endpwent()* functions may fail if:

9184 [EIO] An I/O error has occurred.

9185 In addition, *getpwent()* and *setpwent()* may fail if:

9186 [EMFILE] {OPEN\_MAX} file descriptors are currently open in the calling process.

9187 [ENFILE] The maximum allowable number of files is currently open in the system.

9188 The return value may point to a static area which is overwritten by a subsequent call to |  
 9189 *getpwuid()*, *getpwnam()*, or *getpwent()*.



9190 **EXAMPLES**9191 **Searching the User Database**

9192 The following example uses the *getpwent()* function to get successive entries in the user  
 9193 database, returning a pointer to a **passwd** structure that contains information about each user.  
 9194 The call to *endpwent()* closes the user database and cleans up.

```
9195 #include <pwd.h>
9196 ...
9197 struct passwd *p;
9198 ...
9199 while ((p = getpwent ()) != NULL) {
9200     ...
9201 }
9202 endpwent();
9203 ...
```

9204 **APPLICATION USAGE**

9205 These functions are provided due to their historical usage. Applications should avoid  
 9206 dependencies on fields in the password database, whether the database is a single file, or where  
 9207 in the file system name space the database resides. Applications should use *getpwuid()*  
 9208 whenever possible because it avoids these dependencies.

9209 **RATIONALE**

9210 None.

9211 **FUTURE DIRECTIONS**

9212 None.

9213 **SEE ALSO**

9214 *endgrent()*, *getlogin()*, *getpwnam()*, *getpwuid()*, the Base Definitions volume of  
 9215 IEEE Std 1003.1-200x, **<pwd.h>**

9216 **CHANGE HISTORY**

9217 First released in Issue 4, Version 2.

9218 **Issue 5**

9219 Moved from X/OPEN UNIX extension to BASE.

9220 Normative text previously in the APPLICATION USAGE section is moved to the RETURN  
 9221 VALUE section.

9222 A note indicating that these functions need not be reentrant is added to the DESCRIPTION.

9223 **Issue 6**

9224 In the DESCRIPTION, the note about reentrancy is expanded to cover thread-safety.

9225 **NAME**

9226 endservent, getservbyname, getservbyport, getservent, setservent — network services database  
9227 functions

9228 **SYNOPSIS**

```
9229 #include <netdb.h>

9230 void endservent(void);
9231 struct servent *getservbyname(const char *name, const char *proto);
9232 struct servent *getservbyport(int port, const char *proto);
9233 struct servent *getservent(void);
9234 void setservent(int stayopen);
```

9235 **DESCRIPTION**

9236 These functions shall retrieve information about network services. This information is  
9237 considered to be stored in a database that can be accessed sequentially or randomly. The  
9238 implementation of this database is unspecified. |

9239 The *setservent()* function shall open a connection to the database, and set the next entry to the  
9240 first entry. If the *stayopen* argument is non-zero, the *net* database shall not be closed after each  
9241 call to the *getservent()* function (either directly, or indirectly through one of the other *getserv\*()*  
9242 functions), and the implementation may maintain an open file descriptor for the database.

9243 The *getservent()* function shall read the next entry of the database, opening and closing a  
9244 connection to the database as necessary.

9245 The *getservbyname()* function shall search the database from the beginning and find the first  
9246 entry for which the service name specified by *name* matches the *s\_name* member and the protocol  
9247 name specified by *proto* matches the *s\_proto* member, opening and closing a connection to the  
9248 database as necessary. If *proto* is a null pointer, any value of the *s\_proto* member shall be  
9249 matched.

9250 The *getservbyport()* function shall search the database from the beginning and find the first entry  
9251 for which the port specified by *port* matches the *s\_port* member and the protocol name specified  
9252 by *proto* matches the *s\_proto* member, opening and closing a connection to the database as  
9253 necessary. If *proto* is a null pointer, any value of the *s\_proto* member shall be matched. The *port*  
9254 argument shall be in network byte order.

9255 The *getservbyname()*, *getservbyport()*, and *getservent()* functions shall each return a pointer to a  
9256 **servent** structure, the members of which shall contain the fields of an entry in the network  
9257 services database.

9258 The *endservent()* function shall close the database, releasing any open file descriptor.

9259 These functions need not be reentrant. A function that is not required to be reentrant is not  
9260 required to be thread-safe.

9261 **RETURN VALUE**

9262 Upon successful completion, *getservbyname()*, *getservbyport()*, and *getservent()* return a pointer to  
9263 a **servent** structure if the requested entry was found, and a null pointer if the end of the database  
9264 was reached or the requested entry was not found. Otherwise, a null pointer is returned.

9265 **ERRORS**

9266 No errors are defined.

9267 **EXAMPLES**

9268       None.

9269 **APPLICATION USAGE**9270       The *port* argument of *getservbyport()* need not be compatible with the port values of all address  
9271       families.9272       The *getservbyname()*, *getservbyport()*, and *getservent()* functions may return pointers to static  
9273       data, which may be overwritten by subsequent calls to any of these functions.9274 **RATIONALE**

9275       None.

9276 **FUTURE DIRECTIONS**

9277       None.

9278 **SEE ALSO**9279       *endhostent()*, *endprotoent()*, *htonl()*, *inet\_addr()*, the Base Definitions volume of  
9280       IEEE Std 1003.1-200x, <**netdb.h**>9281 **CHANGE HISTORY**

9282       First released in Issue 6. Derived from the XNS, Issue 5.2 specification.

## 9283 NAME

9284 endutxent, getutxent, getutxid, getutxline, pututxline, setutxent — user accounting database  
9285 functions

## 9286 SYNOPSIS

```
9287 xSI #include <utmpx.h>
9288
9288 void endutxent(void);
9289 struct utmpx *getutxent(void);
9290 struct utmpx *getutxid(const struct utmpx *id);
9291 struct utmpx *getutxline(const struct utmpx *line);
9292 struct utmpx *pututxline(const struct utmpx *utmpx);
9293 void setutxent(void);
9294
```

## 9295 DESCRIPTION

9296 These functions shall provide access to the user accounting database. |

9297 The *getutxent()* function shall read the next entry from the user accounting database. If the |  
9298 database is not already open, it shall open it. If it reaches the end of the database, it shall fail. |

9299 The *getutxid()* function shall search forward from the current point in the database. If the |  
9300 *ut\_type* value of the **utmpx** structure pointed to by *id* is *BOOT\_TIME*, *OLD\_TIME*, or |  
9301 *NEW\_TIME*, then it shall stop when it finds an entry with a matching *ut\_type* value. If the |  
9302 *ut\_type* value is *INIT\_PROCESS*, *LOGIN\_PROCESS*, *USER\_PROCESS*, or *DEAD\_PROCESS*, |  
9303 then it shall stop when it finds an entry whose type is one of these four and whose *ut\_id* member |  
9304 matches the *ut\_id* member of the **utmpx** structure pointed to by *id*. If the end of the database is |  
9305 reached without a match, *getutxid()* shall fail. |

9306 The *getutxline()* function shall search forward from the current point in the database until it |  
9307 finds an entry of the type *LOGIN\_PROCESS* or *USER\_PROCESS* which also has a *ut\_line* value |  
9308 matching that in the **utmpx** structure pointed to by *line*. If the end of the database is reached |  
9309 without a match, *getutxline()* shall fail. |

9310 The *getutxid()* or *getutxline()* function may cache data. For this reason, to use *getutxline()* to |  
9311 search for multiple occurrences, the application shall zero out the static data after each success, |  
9312 or *getutxline()* may return a pointer to the same **utmpx** structure. |

9313 There is one exception to the rule about clearing the structure before further reads are done. The |  
9314 implicit read done by *pututxline()* (if it finds that it is not already at the correct place in the user |  
9315 accounting database) shall not modify the static structure returned by *getutxent()*, *getutxid()*, or |  
9316 *getutxline()*, if the application has modified this structure and passed the pointer back to |  
9317 *pututxline()*.

9318 For all entries that match a request, the *ut\_type* member indicates the type of the entry. Other |  
9319 members of the entry shall contain meaningful data based on the value of the *ut\_type* member as |  
9320 follows:

9321  
9322  
9323  
9324  
9325  
9326  
9327  
9328  
9329  
9330  
9331

<b>ut_type Member</b>	<b>Other Members with Meaningful Data</b>
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>

9332  
9333  
9334  
9335

An implementation that provides extended security controls may impose implementation-defined restrictions on accessing the user accounting database. In particular, the system may deny the existence of some or all of the user accounting database entries associated with users other than the caller.

9336  
9337  
9338  
9339

If the process has appropriate privileges, the *pututxline()* function shall write out the structure into the user accounting database. It shall use *getutxid()* to search for a record that satisfies the request. If this search succeeds, then the entry shall be replaced. Otherwise, a new entry shall be made at the end of the user accounting database.

9340

The *endutxent()* function shall close the user accounting database.

9341  
9342

The *setutxent()* function shall reset 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.

9343  
9344

These functions need not be reentrant. A function that is not required to be reentrant is not required to be thread-safe.

#### 9345 RETURN VALUE

9346  
9347  
9348

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.

9349  
9350

The return value may point to a static area which is overwritten by a subsequent call to *getutxid()* or *getutxline()*.

9351  
9352  
9353

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.

9354

The *endutxent()* and *setutxent()* functions shall not return a value.

#### 9355 ERRORS

9356  
9357

No errors are defined for the *endutxent()*, *getutxent()*, *getutxid()*, *getutxline()*, and *setutxent()* functions.

9358

The *pututxline()* function may fail if:

9359

[EPERM]           The process does not have appropriate privileges.

9360 **EXAMPLES**

9361 None.

9362 **APPLICATION USAGE**9363 The sizes of the arrays in the structure can be found using the *sizeof* operator.9364 **RATIONALE**

9365 None.

9366 **FUTURE DIRECTIONS**

9367 None.

9368 **SEE ALSO**

9369 The Base Definitions volume of IEEE Std 1003.1-200x, &lt;utmpx.h&gt;

9370 **CHANGE HISTORY**

9371 First released in Issue 4, Version 2.

9372 **Issue 5**

9373 Moved from X/OPEN UNIX extension to BASE.

9374 Normative text previously in the APPLICATION USAGE section is moved to the  
9375 DESCRIPTION.

9376 A note indicating that these functions need not be reentrant is added to the DESCRIPTION.

9377 **Issue 6**

9378 In the DESCRIPTION, the note about reentrancy is expanded to cover thread-safety.

9379 **NAME**

9380        **environ** — array of character pointers to the environment strings

9381 **SYNOPSIS**

9382        extern char \*\*environ;

9383 **DESCRIPTION**

9384        Refer to the Base Definitions volume of IEEE Std 1003.1-200x, Chapter 8, Environment Variables  
9385        and *exec*.

9386 **NAME**9387        **erand48** — generate uniformly distributed pseudo-random numbers9388 **SYNOPSIS**9389 **XSI**        #include <stdlib.h>9390        double erand48(unsigned short *xsubi*[3]);

9391

9392 **DESCRIPTION**9393        Refer to *drand48()*.



9394 **NAME**

9395 erf, erff, erfl — error functions

9396 **SYNOPSIS**

9397 #include &lt;math.h&gt;

9398 double erf(double x);

9399 float erff(float x);

9400 long double erfl(long double x);

9401 **DESCRIPTION**

9402 **CX** The functionality described on this reference page is aligned with the ISO C standard. Any  
 9403 conflict between the requirements described here and the ISO C standard is unintentional. This  
 9404 volume of IEEE Std 1003.1-200x defers to the ISO C standard.

9405 These functions shall compute the error function of their argument *x*, defined as:

$$9406 \frac{2}{\sqrt{\pi}} \int_0^x e^{-t^2} dt$$

9407 An application wishing to check for error situations should set *errno* to zero and call  
 9408 *feclearexcept*(FE\_ALL\_EXCEPT) before calling these functions. On return, if *errno* is non-zero or  
 9409 *fetestexcept*(FE\_INVALID | FE\_DIVBYZERO | FE\_OVERFLOW | FE\_UNDERFLOW) is non-  
 9410 zero, an error has occurred.

9411 **RETURN VALUE**

9412 Upon successful completion, these functions shall return the value of the error function.

9413 **MX** If *x* is NaN, a NaN shall be returned.9414 If *x* is ±0, ±0 shall be returned.9415 If *x* is ±Inf, ±1 shall be returned.9416 If *x* is subnormal, a range error may occur, and  $2 * x / \text{sqrt}(\pi)$  should be returned.9417 **ERRORS**

9418 These functions may fail if:

9419 **MX** **Range Error** The result underflows.

9420 If the integer expression (math\_errhandling & MATH\_ERRNO) is non-zero,  
 9421 then *errno* shall be set to [ERANGE]. If the integer expression  
 9422 (math\_errhandling & MATH\_ERREXCEPT) is non-zero, then the underflow  
 9423 floating-point exception shall be raised.

9424 **EXAMPLES**

9425 None.

9426 **APPLICATION USAGE**9427 Underflow occurs when  $|x| < \text{DBL\_MIN} * (\text{sqrt}(\pi)/2)$ .

9428 On error, the expressions (math\_errhandling & MATH\_ERRNO) and (math\_errhandling &  
 9429 MATH\_ERREXCEPT) are independent of each other, but at least one of them must be non-zero.

9430 **RATIONALE**

9431 None.

9432 **FUTURE DIRECTIONS**

9433 None.

9434 **SEE ALSO**9435 *erfc()*, *feclearexcept()*, *fetestexcept()*, *isnan()*, the Base Definitions volume of IEEE Std 1003.1-200x, |  
9436 Section 4.18, Treatment of Error Conditions for Mathematical Functions, <math.h> |9437 **CHANGE HISTORY**

9438 First released in Issue 1. Derived from Issue 1 of the SVID.

9439 **Issue 5**9440 The DESCRIPTION is updated to indicate how an application should check for an error. This  
9441 text was previously published in the APPLICATION USAGE section.9442 **Issue 6**9443 The *erf()* function is no longer marked as an extension.9444 The *erfc()* function is now split out onto its own reference page.9445 The *erff()* and *erfl()* functions are added for alignment with the ISO/IEC 9899:1999 standard.9446 The DESCRIPTION, RETURN VALUE, ERRORS, and APPLICATION USAGE sections are  
9447 revised to align with the ISO/IEC 9899:1999 standard.9448 IEC 60559:1989 standard floating-point extensions over the ISO/IEC 9899:1999 standard are  
9449 marked.

9450 **NAME**

9451           erfc, erfcf, erfcl — complementary error functions

9452 **SYNOPSIS**

9453           #include &lt;math.h&gt;

9454           double erfc(double x);

9455           float erfcf(float x);

9456           long double erfcl(long double x);

9457 **DESCRIPTION**

9458 **CX**       The functionality described on this reference page is aligned with the ISO C standard. Any  
 9459 conflict between the requirements described here and the ISO C standard is unintentional. This  
 9460 volume of IEEE Std 1003.1-200x defers to the ISO C standard.

9461           These functions shall compute the complementary error function  $1.0 - \operatorname{erf}(x)$ .

9462           An application wishing to check for error situations should set *errno* to zero and call  
 9463 *feclearexcept*(FE\_ALL\_EXCEPT) before calling these functions. On return, if *errno* is non-zero or  
 9464 *fetestexcept*(FE\_INVALID | FE\_DIVBYZERO | FE\_OVERFLOW | FE\_UNDERFLOW) is non-  
 9465 zero, an error has occurred.

9466 **RETURN VALUE**9467           Upon successful completion, these functions shall return the value of the complementary error  
9468 function.9469           If the correct value would cause underflow and is not representable, a range error may occur  
9470 **MX**       and either 0.0 (if representable), or an implementation-defined value shall be returned.9471 **MX**       If *x* is NaN, a NaN shall be returned.9472           If *x* is  $\pm 0$ , +1 shall be returned.9473           If *x* is  $-\operatorname{Inf}$ , +2 shall be returned.9474           If *x* is  $+\operatorname{Inf}$ , +0 shall be returned.9475           If the correct value would cause underflow and is representable, a range error may occur and the  
9476 correct value shall be returned.9477 **ERRORS**

9478           These functions may fail if:

9479           Range Error       The result underflows.

9480           If the integer expression (math\_errhandling & MATH\_ERRNO) is non-zero, |  
 9481 then *errno* shall be set to [ERANGE]. If the integer expression |  
 9482 (math\_errhandling & MATH\_ERREXCEPT) is non-zero, then the underflow |  
 9483 floating-point exception shall be raised. |

9484 **EXAMPLES**

9485           None.

9486 **APPLICATION USAGE**9487           The *erfc()* function is provided because of the extreme loss of relative accuracy if *erf(x)* is called  
9488 for large *x* and the result subtracted from 1.0.9489           Note for IEEE Std 754-1985 **double**,  $26.55 < x$  implies *erfc(x)* has underflowed.9490           On error, the expressions (math\_errhandling & MATH\_ERRNO) and (math\_errhandling &  
9491 MATH\_ERREXCEPT) are independent of each other, but at least one of them must be non-zero.

9492 **RATIONALE**

9493 None.

9494 **FUTURE DIRECTIONS**

9495 None.

9496 **SEE ALSO**9497 *erf()*, *feclearexcept()*, *fetetestexcept()*, *isnan()*, the Base Definitions volume of IEEE Std 1003.1-200x, |

9498 Section 4.18, Treatment of Error Conditions for Mathematical Functions, &lt;math.h&gt; |

9499 **CHANGE HISTORY**

9500 First released in Issue 1. Derived from Issue 1 of the SVID.

9501 **Issue 5**9502 The DESCRIPTION is updated to indicate how an application should check for an error. This  
9503 text was previously published in the APPLICATION USAGE section.9504 **Issue 6**9505 The *erfc()* function is no longer marked as an extension.9506 These functions are split out from the *erf()* reference page.9507 The DESCRIPTION, RETURN VALUE, ERRORS, and APPLICATION USAGE sections are  
9508 revised to align with the ISO/IEC 9899:1999 standard.9509 IEC 60559:1989 standard floating-point extensions over the ISO/IEC 9899:1999 standard are  
9510 marked.

9511 **NAME**

9512       erff, erfl — error functions

9513 **SYNOPSIS**

9514       #include &lt;math.h&gt;

9515       float erff(float x);

9516       long double erfl(long double x);

9517 **DESCRIPTION**9518       Refer to *erf()*.

9519 **NAME**9520 `errno` — error return value9521 **SYNOPSIS**9522 `#include <errno.h>`9523 **DESCRIPTION**9524 The lvalue *errno* is used by many functions to return error values. |

9525 Many functions provide an error number in *errno*. It has type **int** and is defined in `<errno.h>`. |  
9526 The value of *errno* shall be defined only after a call to a function for which it is explicitly stated to |  
9527 be set and until it is changed by the next function call or if the application assigns it a value. The |  
9528 value of *errno* should only be examined when it is indicated to be valid by a function's return |  
9529 value. Applications shall obtain the definition of *errno* by the inclusion of `<errno.h>`. No |  
9530 function in this volume of IEEE Std 1003.1-200x shall set *errno* to 0.

9531 It is unspecified whether *errno* is a macro or an identifier declared with external linkage. If a |  
9532 macro definition is suppressed in order to access an actual object, or a program defines an |  
9533 identifier with the name *errno*, the behavior is undefined.

9534 The symbolic values stored in *errno* are documented in the ERRORS sections on all relevant |  
9535 pages.

9536 **RETURN VALUE**

9537 None.

9538 **ERRORS**

9539 None.

9540 **EXAMPLES**

9541 None.

9542 **APPLICATION USAGE**

9543 Previously both POSIX and X/Open documents were more restrictive than the ISO C standard |  
9544 in that they required *errno* to be defined as an external variable, whereas the ISO C standard |  
9545 required only that *errno* be defined as a modifiable lvalue with type **int**. |

9546 A program that uses *errno* for error checking should set it to 0 before a function call, then inspect |  
9547 it before a subsequent function call.

9548 **RATIONALE**

9549 None.

9550 **FUTURE DIRECTIONS**

9551 None.

9552 **SEE ALSO**9553 Section 2.3, the Base Definitions volume of IEEE Std 1003.1-200x, `<errno.h>`9554 **CHANGE HISTORY**

9555 First released in Issue 1. Derived from Issue 1 of the SVID.

9556 **Issue 5**

9557 The following sentence is deleted from the DESCRIPTION: “The value of *errno* is 0 at program |  
9558 start-up, but is never set to 0 by any XSI function”. The DESCRIPTION also no longer states that |  
9559 conforming implementations may support the declaration:

9560 `extern int errno;`

9561 **Issue 6**

9562        Obsolescent text regarding defining *errno* as:

9563        extern int errno

9564        is removed.

9565        Text regarding no function setting *errno* to zero to indicate an error is changed to no function  
9566        shall set *errno* to zero. This is for alignment with the ISO/IEC 9899:1999 standard.

9567 **NAME**

9568 environ, execl, execv, execl, execve, execlp, execvp — execute a file

9569 **SYNOPSIS**

```

9570 #include <unistd.h>

9571 extern char **environ;
9572 int execl(const char *path, const char *arg0, ... /*, (char *)0 */);
9573 int execv(const char *path, char *const argv[]);
9574 int execl(const char *path, const char *arg0, ... /*,
9575           (char *)0, char *const envp[] */);
9576 int execve(const char *path, char *const argv[], char *const envp[]);
9577 int execlp(const char *file, const char *arg0, ... /*, (char *)0 */);
9578 int execvp(const char *file, char *const argv[]);

```

9579 **DESCRIPTION**

9580 The *exec* family of functions shall replace the current process image with a new process image. |  
 9581 The new image shall be constructed from a regular, executable file called the *new process image* |  
 9582 *file*. There shall be no return from a successful *exec*, because the calling process image is overlaid |  
 9583 by the new process image.

9584 When a C-language program is executed as a result of this call, it shall be entered as a C-  
 9585 language function call as follows:

```
9586 int main (int argc, char *argv[]);
```

9587 where *argc* is the argument count and *argv* is an array of character pointers to the arguments  
 9588 themselves. In addition, the following variable:

```
9589 extern char **environ;
```

9590 is initialized as a pointer to an array of character pointers to the environment strings. The *argv*  
 9591 and *environ* arrays are each terminated by a null pointer. The null pointer terminating the *argv*  
 9592 array is not counted in *argc*.

9593 **THR** Conforming multi-threaded applications shall not use the *environ* variable to access or modify  
 9594 any environment variable while any other thread is concurrently modifying any environment  
 9595 variable. A call to any function dependent on any environment variable shall be considered a use  
 9596 of the *environ* variable to access that environment variable.

9597 The arguments specified by a program with one of the *exec* functions shall be passed on to the  
 9598 new process image in the corresponding *main()* arguments.

9599 The argument *path* points to a pathname that identifies the new process image file. |

9600 The argument *file* is used to construct a pathname that identifies the new process image file. If |  
 9601 the *file* argument contains a slash character, the *file* argument shall be used as the pathname for |  
 9602 this file. Otherwise, the path prefix for this file is obtained by a search of the directories passed |  
 9603 as the environment variable *PATH* (see the Base Definitions volume of IEEE Std 1003.1-200x,  
 9604 Chapter 8, Environment Variables). If this environment variable is not present, the results of the  
 9605 search are implementation-defined.

9606 If the process image file is not a valid executable object, and the system does not recognize it as  
 9607 something that cannot be executed (and thus returns [EINVAL]), *execlp()* and *execvp()* shall use  
 9608 the contents of that file as standard input to a command interpreter conforming to *system()*. In  
 9609 this case, the command interpreter becomes the new process image.

9610 The arguments represented by *arg0*,... are pointers to null-terminated character strings. These  
 9611 strings shall constitute the argument list available to the new process image. The list is |



9612 terminated by a null pointer. The argument *arg0* should point to a filename that is associated |  
9613 with the process being started by one of the *exec* functions.

9614 The argument *argv* is an array of character pointers to null-terminated strings. The application |  
9615 shall ensure that the last member of this array is a null pointer. These strings shall constitute the |  
9616 argument list available to the new process image. The value in *argv*[0] should point to a filename |  
9617 that is associated with the process being started by one of the *exec* functions.

9618 The argument *envp* is an array of character pointers to null-terminated strings. These strings |  
9619 shall constitute the environment for the new process image. The *envp* array is terminated by a |  
9620 null pointer.

9621 For those forms not containing an *envp* pointer (*execl*(), *execv*(), *execlp*(), and *execvp*()), the |  
9622 environment for the new process image shall be taken from the external variable *environ* in the |  
9623 calling process.

9624 The number of bytes available for the new process' combined argument and environment lists is |  
9625 {ARG\_MAX}. It is implementation-defined whether null terminators, pointers, and/or any |  
9626 alignment bytes are included in this total.

9627 File descriptors open in the calling process image shall remain open in the new process image, |  
9628 except for those whose close-on-exec flag FD\_CLOEXEC is set. For those file descriptors that |  
9629 remain open, all attributes of the open file description remain unchanged. For any file descriptor |  
9630 that is closed for this reason, file locks are removed as a result of the close as described in *close*(). |  
9631 Locks that are not removed by closing of file descriptors remain unchanged.

9632 Directory streams open in the calling process image shall be closed in the new process image.

9633 The state of the floating-point environment in the new process image shall be set to the default. |

9634 XSI The state of conversion descriptors and message catalog descriptors in the new process image is |  
9635 undefined. For the new process image, the equivalent of:

```
9636 setlocale(LC_ALL, "C")
```

9637 shall be executed at start-up. |

9638 Signals set to the default action (SIG\_DFL) in the calling process image shall be set to the default |  
9639 action in the new process image. Except for SIGCHLD, signals set to be ignored (SIG\_IGN) by |  
9640 the calling process image shall be set to be ignored by the new process image. Signals set to be |  
9641 caught by the calling process image shall be set to the default action in the new process image |  
9642 (see <signal.h>). If the SIGCHLD signal is set to be ignored by the calling process image, it is |  
9643 unspecified whether the SIGCHLD signal is set to be ignored or to the default action in the new |  
9644 XSI process image. After a successful call to any of the *exec* functions, alternate signal stacks are not |  
9645 preserved and the SA\_ONSTACK flag shall be cleared for all signals.

9646 After a successful call to any of the *exec* functions, any functions previously registered by *atexit*() |  
9647 are no longer registered.

9648 XSI If the ST\_NOSUID bit is set for the file system containing the new process image file, then the |  
9649 effective user ID, effective group ID, saved set-user-ID, and saved set-group-ID are unchanged |  
9650 in the new process image. Otherwise, if the set-user-ID mode bit of the new process image file is |  
9651 set, the effective user ID of the new process image shall be set to the user ID of the new process |  
9652 image file. Similarly, if the set-group-ID mode bit of the new process image file is set, the |  
9653 effective group ID of the new process image shall be set to the group ID of the new process |  
9654 image file. The real user ID, real group ID, and supplementary group IDs of the new process |  
9655 image shall remain the same as those of the calling process image. The effective user ID and |  
9656 effective group ID of the new process image shall be saved (as the saved set-user-ID and the |  
9657 saved set-group-ID) for use by *setuid*().

9658	XSI	Any shared memory segments attached to the calling process image shall not be attached to the new process image.
9659		
9660	SEM	Any named semaphores open in the calling process shall be closed as if by appropriate calls to <i>sem_close()</i> .
9661		
9662	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.
9663		
9664	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.
9665		
9666		
9667		
9668		
9669	MF SHM	Memory mappings created in the process are unmapped before the address space is rebuilt for the new process image.
9670		
9671	PS	For the SCHED_FIFO and SCHED_RR 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.
9672		
9673		
9674	TMR	Per-process timers created by the calling process shall be deleted before replacing the current process image with the new process image.
9675		
9676	MSG	All open message queue descriptors in the calling process shall be closed, as described in <i>mq_close()</i> .
9677		
9678	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.
9679		
9680		
9681		
9682		
9683		
9684		
9685	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.
9686		
9687		
9688		
9689	TCT	The initial value of the CPU-time clock of the initial thread of the new process image shall be set to zero.
9690		
9691	TRC	If the calling process is being traced, the new process image shall continue 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.
9692		
9693		
9694		
9695		
9696		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.
9697		
9698		The new process shall inherit at least the following attributes from the calling process image:
9699	XSI	<ul style="list-style-type: none"> <li>• Nice value (see <i>nice()</i>)</li> </ul>
9700	XSI	<ul style="list-style-type: none"> <li>• <i>semadj</i> values (see <i>semop()</i>)</li> </ul>

- 9701           • Process ID
  - 9702           • Parent process ID
  - 9703           • Process group ID
  - 9704           • Session membership
  - 9705           • Real user ID
  - 9706           • Real group ID
  - 9707           • Supplementary group IDs
  - 9708           • Time left until an alarm clock signal (see *alarm()*)
  - 9709           • Current working directory
  - 9710           • Root directory
  - 9711           • File mode creation mask (see *umask()*)
  - 9712 XSI       • File size limit (see *ulimit()*)
  - 9713           • Process signal mask (see *sigprocmask()*)
  - 9714           • Pending signal (see *sigpending()*)
  - 9715           • *tms\_utime*, *tms\_stime*, *tms\_cutime*, and *tms\_cstime* (see *times()*)
  - 9716 XSI       • Resource limits
  - 9717 XSI       • Controlling terminal
  - 9718 XSI       • Interval timers
- 9719           All other process attributes defined in this volume of IEEE Std 1003.1-200x shall be the same in  
 9720           the new and old process images. The inheritance of process attributes not defined by this  
 9721           volume of IEEE Std 1003.1-200x is implementation-defined.
- 9722           A call to any *exec* function from a process with more than one thread shall result in all threads  
 9723           being terminated and the new executable image being loaded and executed. No destructor  
 9724           functions shall be called.
- 9725           Upon successful completion, the *exec* functions shall mark for update the *st\_atime* field of the file.  
 9726           If an *exec* function failed but was able to locate the *process image file*, whether the *st\_atime* field is  
 9727           marked for update is unspecified. Should the *exec* function succeed, the process image file shall  
 9728           be considered to have been opened with *open()*. The corresponding *close()* shall be considered  
 9729           to occur at a time after this open, but before process termination or successful completion of a  
 9730           subsequent call to one of the *exec* functions, *posix\_spawn()*, or *posix\_spawnp()*. The *argv[]* and  
 9731           *envp[]* arrays of pointers and the strings to which those arrays point shall not be modified by a  
 9732           call to one of the *exec* functions, except as a consequence of replacing the process image.
- 9733 XSI       The saved resource limits in the new process image are set to be a copy of the process'  
 9734           corresponding hard and soft limits.
- 9735   **RETURN VALUE**
- 9736           If one of the *exec* functions returns to the calling process image, an error has occurred; the return  
 9737           value shall be  $-1$ , and *errno* shall be set to indicate the error.

9738 **ERRORS**9739 The *exec* functions shall fail if:

9740 [E2BIG] The number of bytes used by the new process image's argument list and  
 9741 environment list is greater than the system-imposed limit of {ARG\_MAX}  
 9742 bytes.

9743 [EACCES] Search permission is denied for a directory listed in the new process image  
 9744 file's path prefix, or the new process image file denies execution permission,  
 9745 or the new process image file is not a regular file and the implementation does  
 9746 not support execution of files of its type.

9747 [EINVAL] The new process image file has the appropriate permission and has a  
 9748 recognized executable binary format, but the system does not support  
 9749 execution of a file with this format.

9750 [ELOOP] A loop exists in symbolic links encountered during resolution of the *path* or *file*  
 9751 argument.

9752 [ENAMETOOLONG]  
 9753 The length of the *path* or *file* arguments exceeds {PATH\_MAX} or a pathname  
 9754 component is longer than {NAME\_MAX}.

9755 [ENOENT] A component of *path* or *file* does not name an existing file or *path* or *file* is an  
 9756 empty string.

9757 [ENOTDIR] A component of the new process image file's path prefix is not a directory.

9758 The *exec* functions, except for *execlp()* and *execvp()*, shall fail if:

9759 [ENOEXEC] The new process image file has the appropriate access permission but has an  
 9760 unrecognized format.

9761 The *exec* functions may fail if:

9762 [ELOOP] More than {SYMLOOP\_MAX} symbolic links were encountered during  
 9763 resolution of the *path* or *file* argument.

9764 [ENAMETOOLONG]  
 9765 As a result of encountering a symbolic link in resolution of the *path* argument,  
 9766 the length of the substituted pathname string exceeded {PATH\_MAX}.

9767 [ENOMEM] The new process image requires more memory than is allowed by the  
 9768 hardware or system-imposed memory management constraints.

9769 [ETXTBSY] The new process image file is a pure procedure (shared text) file that is  
 9770 currently open for writing by some process.

9771 **EXAMPLES**9772 **Using *execl()***

9773 The following example executes the *ls* command, specifying the pathname of the executable  
 9774 (*/bin/ls*) and using arguments supplied directly to the command to produce single-column  
 9775 output.

```
9776 #include <unistd.h>
9777 int ret;
9778 ...
9779 ret = execl ("/bin/ls", "ls", "-l", (char *)0);
```

**9780 Using execl()**

9781 The following example is similar to **Using execl()** (on page 758). In addition, it specifies the  
9782 environment for the new process image using the *env* argument.

```
9783 #include <unistd.h>
9784 int ret;
9785 char *env[] = { "HOME=/usr/home", "LOGNAME=home", (char *)0 };
9786 ...
9787 ret = execl ("/bin/ls", "ls", "-l", (char *)0, env);
```

**9788 Using execlp()**

9789 The following example searches for the location of the *ls* command among the directories  
9790 specified by the *PATH* environment variable.

```
9791 #include <unistd.h>
9792 int ret;
9793 ...
9794 ret = execlp ("ls", "ls", "-l", (char *)0);
```

**9795 Using execv()**

9796 The following example passes arguments to the *ls* command in the *cmd* array.

```
9797 #include <unistd.h>
9798 int ret;
9799 char *cmd[] = { "ls", "-l", (char *)0 };
9800 ...
9801 ret = execv ("/bin/ls", cmd);
```

**9802 Using execve()**

9803 The following example passes arguments to the *ls* command in the *cmd* array, and specifies the  
9804 environment for the new process image using the *env* argument.

```
9805 #include <unistd.h>
9806 int ret;
9807 char *cmd[] = { "ls", "-l", (char *)0 };
9808 char *env[] = { "HOME=/usr/home", "LOGNAME=home", (char *)0 };
9809 ...
9810 ret = execve ("/bin/ls", cmd, env);
```

**9811 Using execvp()**

9812 The following example searches for the location of the *ls* command among the directories  
9813 specified by the *PATH* environment variable, and passes arguments to the *ls* command in the  
9814 *cmd* array.

```
9815 #include <unistd.h>
9816 int ret;
9817 char *cmd[] = { "ls", "-l", (char *)0 };
9818 ...
```

9819           ret = execvp ("ls", cmd);

#### 9820 APPLICATION USAGE

9821           As the state of conversion descriptors and message catalog descriptors in the new process image |  
9822           is undefined, conforming applications should not rely on their use and should close them prior |  
9823           to calling one of the *exec* functions.

9824           Applications that require other than the default POSIX locale should call *setlocale()* with the  
9825           appropriate parameters to establish the locale of the new process.

9826           The *environ* array should not be accessed directly by the application.

#### 9827 RATIONALE

9828           Early proposals required that the value of *argc* passed to *main()* be “one or greater”. This was  
9829           driven by the same requirement in drafts of the ISO C standard. In fact, historical  
9830           implementations have passed a value of zero when no arguments are supplied to the caller of  
9831           the *exec* functions. This requirement was removed from the ISO C standard and subsequently  
9832           removed from this volume of IEEE Std 1003.1-200x as well. The wording, in particular the use of  
9833           the word *should*, requires a Strictly Conforming POSIX Application to pass at least one argument  
9834           to the *exec* function, thus guaranteeing that *argc* be one or greater when invoked by such an  
9835           application. In fact, this is good practice, since many existing applications reference *argv[0]*  
9836           without first checking the value of *argc*.

9837           The requirement on a Strictly Conforming POSIX Application also states that the value passed  
9838           as the first argument be a filename associated with the process being started. Although some  
9839           existing applications pass a pathname rather than a filename in some circumstances, a filename |  
9840           is more generally useful, since the common usage of *argv[0]* is in printing diagnostics. In some  
9841           cases the filename passed is not the actual filename of the file; for example, many  
9842           implementations of the *login* utility use a convention of prefixing a hyphen (‘-’) to the actual  
9843           filename, which indicates to the command interpreter being invoked that it is a “login shell”.

9844           Some implementations can *exec* shell scripts. |

9845           One common historical implementation is that the *execl()*, *execv()*, *execle()*, and *execve()*  
9846           functions return an [ENOEXEC] error for any file not recognizable as executable, including a  
9847           shell script. When the *execlp()* and *execvp()* functions encounter such a file, they assume the file  
9848           to be a shell script and invoke a known command interpreter to interpret such files. These  
9849           implementations of *execvp()* and *execlp()* only give the [ENOEXEC] error in the rare case of a  
9850           problem with the command interpreter’s executable file. Because of these implementations, the  
9851           [ENOEXEC] error is not mentioned for *execlp()* or *execvp()*, although implementations can still  
9852           give it.

9853           Another way that some historical implementations handle shell scripts is by recognizing the first  
9854           two bytes of the file as the character string “#!” and using the remainder of the first line of the  
9855           file as the name of the command interpreter to execute.

9856           Some implementations provide a third argument to *main()* called *envp*. This is defined as a  
9857           pointer to the environment. The ISO C standard specifies invoking *main()* with two arguments,  
9858           so implementations must support applications written this way. Since this volume of  
9859           IEEE Std 1003.1-200x defines the global variable *environ*, which is also provided by historical  
9860           implementations and can be used anywhere that *envp* could be used, there is no functional need  
9861           for the *envp* argument. Applications should use the *getenv()* function rather than accessing the  
9862           environment directly via either *envp* or *environ*. Implementations are required to support the  
9863           two-argument calling sequence, but this does not prohibit an implementation from supporting  
9864           *envp* as an optional third argument.

9865 This volume of IEEE Std 1003.1-200x specifies that signals set to SIG\_IGN remain set to  
 9866 SIG\_IGN, and that the process signal mask be unchanged across an *exec*. This is consistent with  
 9867 historical implementations, and it permits some useful functionality, such as the *nohup*  
 9868 command. However, it should be noted that many existing applications wrongly assume that  
 9869 they start with certain signals set to the default action and/or unblocked. In particular,  
 9870 applications written with a simpler signal model that does not include blocking of signals, such  
 9871 as the one in the ISO C standard, may not behave properly if invoked with some signals blocked.  
 9872 Therefore, it is best not to block or ignore signals across *execs* without explicit reason to do so,  
 9873 and especially not to block signals across *execs* of arbitrary (not closely co-operating) programs.

9874 The *exec* functions always save the value of the effective user ID and effective group ID of the  
 9875 process at the completion of the *exec*, whether or not the set-user-ID or the set-group-ID bit of  
 9876 the process image file is set.

9877 The statement about *argv*[] and *envp*[] being constants is included to make explicit to future  
 9878 writers of language bindings that these objects are completely constant. Due to a limitation of  
 9879 the ISO C standard, it is not possible to state that idea in standard C. Specifying two levels of  
 9880 *const-qualification* for the *argv*[] and *envp*[] parameters for the *exec* functions may seem to be the  
 9881 natural choice, given that these functions do not modify either the array of pointers or the  
 9882 characters to which the function points, but this would disallow existing correct code. Instead,  
 9883 only the array of pointers is noted as constant. The table of assignment compatibility for *dst=src*,  
 9884 derived from the ISO C standard summarizes the compatibility:

9885	<i>dst:</i>	char *[]	const char *[]	char *const[]	const char *const[]
9886	<i>src:</i>				
9887	char *[]	VALID	—	VALID	—
9888	const char *[]	—	VALID	—	VALID
9889	char * const []	—	—	VALID	—
9890	const char *const[]	—	—	—	VALID

9891 Since all existing code has a source type matching the first row, the column that gives the most  
 9892 valid combinations is the third column. The only other possibility is the fourth column, but  
 9893 using it would require a cast on the *argv* or *envp* arguments. It is unfortunate that the fourth  
 9894 column cannot be used, because the declaration a non-expert would naturally use would be that  
 9895 in the second row.

9896 The ISO C standard and this volume of IEEE Std 1003.1-200x do not conflict on the use of  
 9897 *environ*, but some historical implementations of *environ* may cause a conflict. As long as *environ*  
 9898 is treated in the same way as an entry point (for example, *fork*()), it conforms to both standards.  
 9899 A library can contain *fork*(), but if there is a user-provided *fork*(), that *fork*() is given precedence  
 9900 and no problem ensues. The situation is similar for *environ*: the definition in this volume of  
 9901 IEEE Std 1003.1-200x is to be used if there is no user-provided *environ* to take precedence. At  
 9902 least three implementations are known to exist that solve this problem.

9903 [E2BIG] The limit {ARG\_MAX} applies not just to the size of the argument list, but to  
 9904 the sum of that and the size of the environment list.

9905 [EFAULT] Some historical systems return [EFAULT] rather than [ENOEXEC] when the  
 9906 new process image file is corrupted. They are non-conforming.

9907 [EINVAL] This error condition was added to IEEE Std 1003.1-200x to allow an  
 9908 implementation to detect executable files generated for different architectures,  
 9909 and indicate this situation to the application. Historical implementations of  
 9910 shells, *execvp*(), and *execlp*() that encounter an [ENOEXEC] error will execute  
 9911 a shell on the assumption that the file is a shell script. This will not produce  
 9912 the desired effect when the file is a valid executable for a different

9913 architecture. An implementation may now choose to avoid this problem by  
 9914 returning [EINVAL] when a valid executable for a different architecture is  
 9915 encountered. Some historical implementations return [EINVAL] to indicate  
 9916 that the *path* argument contains a character with the high order bit set. The  
 9917 standard developers chose to deviate from historical practice for the following  
 9918 reasons:

- 9919 1. The new utilization of [EINVAL] will provide some measure of utility to  
 9920 the user community.
- 9921 2. Historical use of [EINVAL] is not acceptable in an internationalized  
 9922 operating environment.

9923 [ENAMETOOLONG]

9924 Since the file pathname may be constructed by taking elements in the *PATH*  
 9925 variable and putting them together with the filename, the  
 9926 [ENAMETOOLONG] error condition could also be reached this way.

9927 [ETXTBSY]

9928 System V returns this error when the executable file is currently open for  
 9929 writing by some process. This volume of IEEE Std 1003.1-200x neither requires  
 nor prohibits this behavior.

9930 Other systems (such as System V) may return [EINTR] from *exec*. This is not addressed by this  
 9931 volume of IEEE Std 1003.1-200x, but implementations may have a window between the call to  
 9932 *exec* and the time that a signal could cause one of the *exec* calls to return with [EINTR].

9933 An explicit statement regarding the floating-point environment (as defined in the `<fenv.h>`  
 9934 header) was added to make it clear that the floating-point environment is set to its default when  
 9935 a call to one of the *exec* functions succeeds. The requirements for inheritance or setting to the  
 9936 default for other process and thread start-up functions is covered by more generic statements in  
 9937 their descriptions and can be summarized as follows:

9938 <i>posix_spawn()</i>	Set to default.
9939 <i>fork()</i>	Inherit.
9940 <i>pthread_create()</i>	Inherit.

#### 9941 FUTURE DIRECTIONS

9942 None.

#### 9943 SEE ALSO

9944 *alarm()*, *atexit()*, *chmod()*, *close()*, *exit()*, *fcntl()*, *fork()*, *fstatvfs()*, *getenv()*, *getitimer()*, *getrlimit()*,  
 9945 *mmap()*, *nice()*, *posix\_spawn()*, *posix\_trace\_eventid\_open()*, *posix\_trace\_shutdown()*,  
 9946 *posix\_trace\_trid\_eventid\_open()*, *putenv()*, *semop()*, *setlocale()*, *shmat()*, *sigaction()*, *sigaltstack()*,  
 9947 *sigpending()*, *sigprocmask()*, *system()*, *times()*, *ulimit()*, *umask()*, the Base Definitions volume of  
 9948 IEEE Std 1003.1-200x, `<unistd.h>`, the Base Definitions volume of IEEE Std 1003.1-200x, Chapter  
 9949 11, General Terminal Interface

#### 9950 CHANGE HISTORY

9951 First released in Issue 1. Derived from Issue 1 of the SVID.

#### 9952 Issue 5

9953 The DESCRIPTION is updated for alignment with the POSIX Realtime Extension and the POSIX  
 9954 Threads Extension.

9955 Large File Summit extensions are added.



9956 **Issue 6**

9957 The following new requirements on POSIX implementations derive from alignment with the |  
9958 Single UNIX Specification:

9959 • In the DESCRIPTION, behavior is defined for when the process image file is not a valid |  
9960 executable.

9961 • In this issue, `_POSIX_SAVED_IDS` is mandated, thus the effective user ID and effective group |  
9962 ID of the new process image shall be saved (as the saved set-user-ID and the saved set- |  
9963 group-ID) for use by the `setuid()` function.

9964 • The [ELOOP] mandatory error condition is added.

9965 • A second [ENAMETOOLONG] is added as an optional error condition.

9966 • The [ETXTBSY] optional error condition is added.

9967 The following changes were made to align with the IEEE P1003.1a draft standard:

9968 • The [EINVAL] mandatory error condition is added.

9969 • The [ELOOP] optional error condition is added.

9970 The description of CPU-time clock semantics is added for alignment with IEEE Std 1003.1d-1999.

9971 The DESCRIPTION is updated for alignment with IEEE Std 1003.1j-2000 by adding semantics for |  
9972 typed memory.

9973 The DESCRIPTION is updated to avoid use of the term “must” for application requirements.

9974 The description of tracing semantics is added for alignment with IEEE Std 1003.1q-2000. |

9975 IEEE PASC Interpretation 1003.1 #132 is applied. |

9976 The DESCRIPTION is updated to make it explicit that the floating-point environment in the new |  
9977 process image is set to the default. |

## 9978 NAME

9979 exit, \_Exit, \_exit — terminate a process

## 9980 SYNOPSIS

9981 #include &lt;stdlib.h&gt;

9982 void exit(int status);

9983 void \_Exit(int status);

9984 #include &lt;unistd.h&gt;

9985 void \_exit(int status);

## 9986 DESCRIPTION

9987 CX The functionality described on this reference page for the *exit()* and *\_Exit()* functions is aligned  
 9988 with the ISO C standard. Any conflict between the requirements described here and the ISO C  
 9989 standard are unintentional. This volume of IEEE Std 1003.1-200x defers to the ISO C standard.

9990 CX The value of *status* may be 0, EXIT\_SUCCESS, EXIT\_FAILURE, or any other value, though only  
 9991 the least significant 8 bits (that is, *status* & 0377) shall be available to a waiting parent process.

9992 The *exit()* function shall first call all functions registered by *atexit()*, in the reverse order of their  
 9993 registration, except that a function is called after any previously registered functions that had  
 9994 already been called at the time it was registered. Each function is called as many times as it was  
 9995 registered. If, during the call to any such function, a call to the *longjmp()* function is made that  
 9996 would terminate the call to the registered function, the behavior is undefined.

9997 If a function registered by a call to *atexit()* fails to return, the remaining registered functions shall  
 9998 not be called and the rest of the *exit()* processing shall not be completed. If *exit()* is called more  
 9999 than once, the behavior is undefined.

10000 The *exit()* function shall then flush all open streams with unwritten buffered data, close all open  
 10001 streams, and remove all files created by *tmpfile()*. Finally, control shall be terminated with the  
 10002 consequences described below.

10003 CX The *\_Exit()* and *\_exit()* functions shall be functionally equivalent.

10004 CX The *\_Exit()* and *\_exit()* functions shall not call functions registered with *atexit()* nor any  
 10005 registered signal handlers. Whether open streams are flushed or closed, or temporary files are  
 10006 removed is implementation-defined. Finally, the calling process is terminated with the  
 10007 consequences described below.

10008 CX These functions shall terminate the calling process with the following consequences:

10009 **Note:** These consequences are all extensions to the ISO C standard and are not further CX shaded.  
 10010 However, XSI extensions are shaded.

10011 XSI • All of the file descriptors, directory streams, conversion descriptors, and message catalog  
 10012 descriptors open in the calling process shall be closed.

10013 XSI • If the parent process of the calling process is executing a *wait()* or *waitpid()*, and has neither  
 10014 set its SA\_NOCLDWAIT flag nor set SIGCHLD to SIG\_IGN, it shall be notified of the calling  
 10015 process' termination and the low-order eight bits (that is, bits 0377) of *status* are made  
 10016 available to it. If the parent is not waiting, the child's status shall be made available to it  
 10017 when the parent subsequently executes *wait()* or *waitpid()*.

10018 XSI The semantics of the *waitid()* function shall be equivalent to *wait()*.

10019 XSI • If the parent process of the calling process is not executing a *wait()* or *waitpid()*, and has  
 10020 neither set its SA\_NOCLDWAIT flag nor set SIGCHLD to SIG\_IGN, the calling process shall  
 10021 be transformed into a *zombie process*. A *zombie process* is an inactive process and it shall be

10022		deleted at some later time when its parent process executes <i>wait()</i> or <i>waitpid()</i> .	
10023	XSI	The semantics of the <i>waitid()</i> function shall be equivalent to <i>wait()</i> .	
10024		• Termination of a process does not directly terminate its children. The sending of a SIGHUP signal as described below indirectly terminates children in some circumstances.	
10025			
10026		• Either:	
10027		If the implementation supports the SIGCHLD signal, a SIGCHLD shall be sent to the parent process.	
10028			
10029		Or:	
10030	XSI	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 is sent to the parent process.	
10031			
10032			
10033			
10034		• The parent process ID of all of the calling process' existing child processes and zombie processes shall be set to the process ID of an implementation-defined system process. That is, these processes shall be inherited by a special system process.	
10035			
10036			
10037	XSI	• Each attached shared-memory segment is detached and the value of <i>shm_nattch</i> (see <i>shmget()</i> ) in the data structure associated with its shared memory ID shall be decremented by 1.	
10038			
10039			
10040	XSI	• For each semaphore for which the calling process has set a <i>semadj</i> value (see <i>semop()</i> ), that value shall be added to the <i>semval</i> of the specified semaphore.	
10041			
10042		• 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.	
10043			
10044		• If the process is a controlling process, the controlling terminal associated with the session shall be disassociated from the session, allowing it to be acquired by a new controlling process.	
10045			
10046			
10047		• 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.	
10048			
10049			
10050	SEM	• All open named semaphores in the calling process shall be closed as if by appropriate calls to <i>sem_close()</i> .	
10051			
10052	ML	• Any memory locks established by the 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 <i>_Exit()</i> or <i>_exit()</i> .	
10053			
10054			
10055			
10056	MF SHM	• Memory mappings created in the process shall be unmapped before the process is destroyed.	
10057			
10058	TYM	• Any blocks of typed memory that were mapped in the calling process shall be unmapped, as if <i>munmap()</i> was implicitly called to unmap them.	
10059			
10060	MSG	• All open message queue descriptors in the calling process shall be closed as if by appropriate calls to <i>mq_close()</i> .	
10061			
10062	AIO	• Any outstanding cancelable asynchronous I/O operations may be canceled. Those asynchronous I/O operations that are not canceled shall complete as if the <i>_Exit()</i> or <i>_exit()</i> operation had not yet occurred, but any associated signal notifications shall be suppressed.	
10063			
10064			

10065 The `_Exit()` or `_exit()` operation may block awaiting such I/O completion. Whether any I/O  
 10066 is canceled, and which I/O may be canceled upon `_Exit()` or `_exit()`, is implementation-  
 10067 defined.

10068 • Threads terminated by a call to `_Exit()` or `_exit()` shall not invoke their cancelation cleanup  
 10069 handlers or per-thread data destructors.

10070 TRC • If the calling process is a trace controller process, any trace streams that were created by the  
 10071 calling process shall be shut down as described by the `posix_trace_shutdown()` function, and  
 10072 any process' mapping of trace event names to trace event type identifiers built for these trace  
 10073 streams may be deallocated.

#### 10074 RETURN VALUE

10075 These functions do not return.

#### 10076 ERRORS

10077 No errors are defined.

#### 10078 EXAMPLES

10079 None.

#### 10080 APPLICATION USAGE

10081 Normally applications should use `exit()` rather than `_Exit()` or `_exit()`.

#### 10082 RATIONALE

##### 10083 Process Termination

10084 Early proposals drew a distinction between normal and abnormal process termination.  
 10085 Abnormal termination was caused only by certain signals and resulted in implementation-  
 10086 defined “actions”, as discussed below. Subsequent proposals distinguished three types of  
 10087 termination: *normal termination* (as in the current specification), *simple abnormal termination*, and  
 10088 *abnormal termination with actions*. Again the distinction between the two types of abnormal  
 10089 termination was that they were caused by different signals and that implementation-defined  
 10090 actions would result in the latter case. Given that these actions were completely  
 10091 implementation-defined, the early proposals were only saying when the actions could occur and  
 10092 how their occurrence could be detected, but not what they were. This was of little or no use to  
 10093 conforming applications, and thus the distinction is not made in this volume of |  
 10094 IEEE Std 1003.1-200x.

10095 The implementation-defined actions usually include, in most historical implementations, the  
 10096 creation of a file named **core** in the current working directory of the process. This file contains an  
 10097 image of the memory of the process, together with descriptive information about the process,  
 10098 perhaps sufficient to reconstruct the state of the process at the receipt of the signal.

10099 There is a potential security problem in creating a **core** file if the process was set-user-ID and the  
 10100 current user is not the owner of the program, if the process was set-group-ID and none of the  
 10101 user's groups match the group of the program, or if the user does not have permission to write in  
 10102 the current directory. In this situation, an implementation either should not create a **core** file or  
 10103 should make it unreadable by the user.

10104 Despite the silence of this volume of IEEE Std 1003.1-200x on this feature, applications are  
 10105 advised not to create files named **core** because of potential conflicts in many implementations.  
 10106 Some historical implementations use a different name than **core** for the file, such as by  
 10107 appending the process ID to the filename.

10108 **Terminating a Process**

10109 It is important that the consequences of process termination as described occur regardless of  
10110 whether the process called `_exit()` (perhaps indirectly through `exit()`) or instead was terminated  
10111 due to a signal or for some other reason. Note that in the specific case of `exit()` this means that  
10112 the *status* argument to `exit()` is treated in the same way as the *status* argument to `_exit()`.

10113 A language other than C may have other termination primitives than the C-language `exit()`  
10114 function, and programs written in such a language should use its native termination primitives,  
10115 but those should have as part of their function the behavior of `_exit()` as described.  
10116 Implementations in languages other than C are outside the scope of the present version of this  
10117 volume of IEEE Std 1003.1-200x, however.

10118 As required by the ISO C standard, using **return** from `main()` has the same behavior (other than  
10119 with respect to language scope issues) as calling `exit()` with the returned value. Reaching the end  
10120 of the `main()` function has the same behavior as calling `exit(0)`.

10121 A value of zero (or `EXIT_SUCCESS`, which is required to be zero) for the argument *status*  
10122 conventionally indicates successful termination. This corresponds to the specification for `exit()`  
10123 in the ISO C standard. The convention is followed by utilities such as *make* and various shells,  
10124 which interpret a zero status from a child process as success. For this reason, applications should  
10125 not call `exit(0)` or `_exit(0)` when they terminate unsuccessfully; for example, in signal-catching  
10126 functions.

10127 Historically, the implementation-defined process that inherits children whose parents have  
10128 terminated without waiting on them is called *init* and has a process ID of 1.

10129 The sending of a `SIGHUP` to the foreground process group when a controlling process  
10130 terminates corresponds to somewhat different historical implementations. In System V, the  
10131 kernel sends a `SIGHUP` on termination of (essentially) a controlling process. In 4.2 BSD, the  
10132 kernel does not send `SIGHUP` in a case like this, but the termination of a controlling process is  
10133 usually noticed by a system daemon, which arranges to send a `SIGHUP` to the foreground  
10134 process group with the `vhangup()` function. However, in 4.2 BSD, due to the behavior of the  
10135 shells that support job control, the controlling process is usually a shell with no other processes  
10136 in its process group. Thus, a change to make `_exit()` behave this way in such systems should not  
10137 cause problems with existing applications.

10138 The termination of a process may cause a process group to become orphaned in either of two  
10139 ways. The connection of a process group to its parent(s) outside of the group depends on both  
10140 the parents and their children. Thus, a process group may be orphaned by the termination of the  
10141 last connecting parent process outside of the group or by the termination of the last direct  
10142 descendant of the parent process(es). In either case, if the termination of a process causes a  
10143 process group to become orphaned, processes within the group are disconnected from their job  
10144 control shell, which no longer has any information on the existence of the process group.  
10145 Stopped processes within the group would languish forever. In order to avoid this problem,  
10146 newly orphaned process groups that contain stopped processes are sent a `SIGHUP` signal and a  
10147 `SIGCONT` signal to indicate that they have been disconnected from their session. The `SIGHUP`  
10148 signal causes the process group members to terminate unless they are catching or ignoring  
10149 `SIGHUP`. Under most circumstances, all of the members of the process group are stopped if any  
10150 of them are stopped.

10151 The action of sending a `SIGHUP` and a `SIGCONT` signal to members of a newly orphaned  
10152 process group is similar to the action of 4.2 BSD, which sends `SIGHUP` and `SIGCONT` to each  
10153 stopped child of an exiting process. If such children exit in response to the `SIGHUP`, any  
10154 additional descendants receive similar treatment at that time. In this volume of  
10155 IEEE Std 1003.1-200x, the signals are sent to the entire process group at the same time. Also, in

10156 this volume of IEEE Std 1003.1-200x, but not in 4.2 BSD, stopped processes may be orphaned,  
10157 but may be members of a process group that is not orphaned; therefore, the action taken at  
10158 `_exit()` must consider processes other than child processes.

10159 It is possible for a process group to be orphaned by a call to `setpgid()` or `setsid()`, as well as by  
10160 process termination. This volume of IEEE Std 1003.1-200x does not require sending `SIGHUP` and  
10161 `SIGCONT` in those cases, because, unlike process termination, those cases are not caused  
10162 accidentally by applications that are unaware of job control. An implementation can choose to  
10163 send `SIGHUP` and `SIGCONT` in those cases as an extension; such an extension must be  
10164 documented as required in `<signal.h>`.

10165 The ISO/IEC 9899:1999 standard adds the `_Exit()` function that results in immediate program  
10166 termination without triggering signals or `atexit()`-registered functions. In IEEE Std 1003.1-200x,  
10167 this is equivalent to the `_exit()` function.

#### 10168 FUTURE DIRECTIONS

10169 None.

#### 10170 SEE ALSO

10171 `atexit()`, `close()`, `fclose()`, `longjmp()`, `posix_trace_shutdown()`, `posix_trace_trid_eventid_open()`,  
10172 `semop()`, `shmget()`, `sigaction()`, `wait()`, `waitid()`, `waitpid()`, the Base Definitions volume of  
10173 IEEE Std 1003.1-200x, `<stdlib.h>`, `<unistd.h>`

#### 10174 CHANGE HISTORY

10175 First released in Issue 1. Derived from Issue 1 of the SVID.

#### 10176 Issue 5

10177 The DESCRIPTION is updated for alignment with the POSIX Realtime Extension and the POSIX  
10178 Threads Extension.

10179 Interactions with the `SA_NOCLDWAIT` flag and `SIGCHLD` signal are further clarified.

10180 The values of `status` from `exit()` are better described.

#### 10181 Issue 6

10182 Extensions beyond the ISO C standard are now marked.

10183 The DESCRIPTION is updated for alignment with IEEE Std 1003.1j-2000 by adding semantics for  
10184 typed memory.

10185 The following changes are made for alignment with the ISO/IEC 9899:1999 standard:

- 10186 • The `_Exit()` function is included.
- 10187 • The DESCRIPTION is updated.

10188 The description of tracing semantics is added for alignment with IEEE Std 1003.1q-2000.

10189 References to the `wait3()` function are removed.

10190 **NAME**

10191 exp, expf, expl — exponential function

10192 **SYNOPSIS**

10193 #include &lt;math.h&gt;

10194 double exp(double x);

10195 float expf(float x);

10196 long double expl(long double x);

10197 **DESCRIPTION**

10198 **CX** The functionality described on this reference page is aligned with the ISO C standard. Any  
 10199 conflict between the requirements described here and the ISO C standard is unintentional. This  
 10200 volume of IEEE Std 1003.1-200x defers to the ISO C standard.

10201 These functions shall compute the base-*e* exponential of *x*.

10202 An application wishing to check for error situations should set *errno* to zero and call  
 10203 *feclearexcept*(FE\_ALL\_EXCEPT) before calling these functions. On return, if *errno* is non-zero or  
 10204 *fetestexcept*(FE\_INVALID | FE\_DIVBYZERO | FE\_OVERFLOW | FE\_UNDERFLOW) is non-  
 10205 zero, an error has occurred.

10206 **RETURN VALUE**10207 Upon successful completion, these functions shall return the exponential value of *x*.

10208 If the correct value would cause overflow, a range error shall occur and *exp()*, *expf()*, and *expl()*  
 10209 shall return the value of the macro HUGE\_VAL, HUGE\_VALF, and HUGE\_VALL, respectively.

10210 If the correct value would cause underflow, and is not representable, a range error may occur,  
 10211 **MX** and either 0.0 (if supported), or an implementation-defined value shall be returned.

10212 **MX** If *x* is NaN, a NaN shall be returned.10213 If *x* is ±0, 1 shall be returned.10214 If *x* is -Inf, +0 shall be returned.10215 If *x* is +Inf, *x* shall be returned.

10216 If the correct value would cause underflow, and is representable, a range error may occur and  
 10217 the correct value shall be returned.

10218 **ERRORS**

10219 These functions shall fail if:

10220 Range Error The result overflows.

10221 If the integer expression (math\_errhandling & MATH\_ERRNO) is non-zero, |  
 10222 then *errno* shall be set to [ERANGE]. If the integer expression |  
 10223 (math\_errhandling & MATH\_ERREXCEPT) is non-zero, then the overflow |  
 10224 floating-point exception shall be raised. |

10225 These functions may fail if:

10226 Range Error The result underflows.

10227 If the integer expression (math\_errhandling & MATH\_ERRNO) is non-zero, |  
 10228 then *errno* shall be set to [ERANGE]. If the integer expression |  
 10229 (math\_errhandling & MATH\_ERREXCEPT) is non-zero, then the underflow |  
 10230 floating-point exception shall be raised. |

10231 **EXAMPLES**

10232 None.

10233 **APPLICATION USAGE**

10234 Note that for IEEE Std 754-1985 **double**,  $709.8 < x$  implies  $\exp(x)$  has overflowed. The value  $x <$   
10235  $-708.4$  implies  $\exp(x)$  has underflowed.

10236 On error, the expressions `(math_errhandling & MATH_ERRNO)` and `(math_errhandling &`  
10237 `MATH_ERREXCEPT)` are independent of each other, but at least one of them must be non-zero.

10238 **RATIONALE**

10239 None.

10240 **FUTURE DIRECTIONS**

10241 None.

10242 **SEE ALSO**

10243 `feclearexcept()`, `fetestexcept()`, `isnan()`, `log()`, the Base Definitions volume of IEEE Std 1003.1-200x, |  
10244 Section 4.18, Treatment of Error Conditions for Mathematical Functions, `<math.h>` |

10245 **CHANGE HISTORY**

10246 First released in Issue 1. Derived from Issue 1 of the SVID.

10247 **Issue 5**

10248 The DESCRIPTION is updated to indicate how an application should check for an error. This  
10249 text was previously published in the APPLICATION USAGE section.

10250 **Issue 6**10251 The `expf()` and `expl()` functions are added for alignment with the ISO/IEC 9899:1999 standard.

10252 The DESCRIPTION, RETURN VALUE, ERRORS, and APPLICATION USAGE sections are  
10253 revised to align with the ISO/IEC 9899:1999 standard.

10254 IEC 60559:1989 standard floating-point extensions over the ISO/IEC 9899:1999 standard are  
10255 marked.



10256 **NAME**

10257 exp2, exp2f, exp2l — exponential base 2 functions

10258 **SYNOPSIS**

10259 #include &lt;math.h&gt;

10260 double exp2(double x);

10261 float exp2f(float x);

10262 long double exp2l(long double x);

10263 **DESCRIPTION**

10264 **CX** The functionality described on this reference page is aligned with the ISO C standard. Any  
 10265 conflict between the requirements described here and the ISO C standard is unintentional. This  
 10266 volume of IEEE Std 1003.1-200x defers to the ISO C standard.

10267 These functions shall compute the base-2 exponential of  $x$ .

10268 An application wishing to check for error situations should set *errno* to zero and call  
 10269 *feclearexcept*(FE\_ALL\_EXCEPT) before calling these functions. On return, if *errno* is non-zero or  
 10270 *fetestexcept*(FE\_INVALID | FE\_DIVBYZERO | FE\_OVERFLOW | FE\_UNDERFLOW) is non-  
 10271 zero, an error has occurred.

10272 **RETURN VALUE**10273 Upon successful completion, these functions shall return  $2^x$ .

10274 If the correct value would cause overflow, a range error shall occur and *exp2()*, *exp2f()*, and  
 10275 *exp2l()* shall return the value of the macro HUGE\_VAL, HUGE\_VALF, and HUGE\_VALL,  
 10276 respectively.

10277 If the correct value would cause underflow, and is not representable, a range error may occur,  
 10278 **MX** and either 0.0 (if supported), or an implementation-defined value shall be returned.

10279 **MX** If  $x$  is NaN, a NaN shall be returned.10280 If  $x$  is  $\pm 0$ , 1 shall be returned.10281 If  $x$  is  $-\text{Inf}$ , +0 shall be returned.10282 If  $x$  is  $+\text{Inf}$ ,  $x$  shall be returned.

10283 If the correct value would cause underflow, and is representable, a range error may occur and  
 10284 the correct value shall be returned.

10285 **ERRORS**

10286 These functions shall fail if:

10287 Range Error The result overflows.

10288 If the integer expression (math\_errhandling & MATH\_ERRNO) is non-zero, |  
 10289 then *errno* shall be set to [ERANGE]. If the integer expression |  
 10290 (math\_errhandling & MATH\_ERREXCEPT) is non-zero, then the overflow |  
 10291 floating-point exception shall be raised. |

10292 These functions may fail if:

10293 Range Error The result underflows.

10294 If the integer expression (math\_errhandling & MATH\_ERRNO) is non-zero, |  
 10295 then *errno* shall be set to [ERANGE]. If the integer expression |  
 10296 (math\_errhandling & MATH\_ERREXCEPT) is non-zero, then the underflow |  
 10297 floating-point exception shall be raised. |

10298 **EXAMPLES**

10299 None.

10300 **APPLICATION USAGE**

10301 For IEEE Std 754-1985 **double**,  $1024 \leq x$  implies  $\text{exp2}(x)$  has overflowed. The value  $x < -1022$   
10302 implies  $\text{exp}(x)$  has underflowed.

10303 On error, the expressions `(math_errhandling & MATH_ERRNO)` and `(math_errhandling &`  
10304 `MATH_ERREXCEPT)` are independent of each other, but at least one of them must be non-zero.

10305 **RATIONALE**

10306 None.

10307 **FUTURE DIRECTIONS**

10308 None.

10309 **SEE ALSO**

10310 *exp()*, *feclearexcept()*, *fetestexcept()*, *isnan()*, *log()*, the Base Definitions volume of |  
10311 IEEE Std 1003.1-200x, Section 4.18, Treatment of Error Conditions for Mathematical Functions, |  
10312 `<math.h>`

10313 **CHANGE HISTORY**

10314 First released in Issue 6. Derived from the ISO/IEC 9899:1999 standard.

10315 **NAME**

10316 expm1, expm1f, expm1l — compute exponential functions

10317 **SYNOPSIS**

10318 #include &lt;math.h&gt;

10319 double expm1(double x);

10320 float expm1f(float x);

10321 long double expm1l(long double x);

10322 **DESCRIPTION**

10323 **CX** The functionality described on this reference page is aligned with the ISO C standard. Any  
 10324 conflict between the requirements described here and the ISO C standard is unintentional. This  
 10325 volume of IEEE Std 1003.1-200x defers to the ISO C standard.

10326 These functions shall compute  $e^x-1.0$ .

10327 An application wishing to check for error situations should set *errno* to zero and call  
 10328 *feclearexcept*(FE\_ALL\_EXCEPT) before calling these functions. On return, if *errno* is non-zero or  
 10329 *fetestexcept*(FE\_INVALID | FE\_DIVBYZERO | FE\_OVERFLOW | FE\_UNDERFLOW) is non-  
 10330 zero, an error has occurred.

10331 **RETURN VALUE**10332 Upon successful completion, these functions return  $e^x-1.0$ .

10333 If the correct value would cause overflow, a range error shall occur and *expm1()*, *expm1f()*, and  
 10334 *expm1l()* shall return the value of the macro HUGE\_VAL, HUGE\_VALF, and HUGE\_VALL,  
 10335 respectively.

10336 **MX** If *x* is NaN, a NaN shall be returned.10337 If *x* is  $\pm 0$ ,  $\pm 0$  shall be returned.10338 If *x* is  $-\text{Inf}$ ,  $-1$  shall be returned.10339 If *x* is  $+\text{Inf}$ , *x* shall be returned.10340 If *x* is subnormal, a range error may occur and *x* should be returned.10341 **ERRORS**

10342 These functions shall fail if:

10343 **Range Error** The result overflows.

10344 If the integer expression (math\_errhandling & MATH\_ERRNO) is non-zero, |  
 10345 then *errno* shall be set to [ERANGE]. If the integer expression |  
 10346 (math\_errhandling & MATH\_ERREXCEPT) is non-zero, then the overflow |  
 10347 floating-point exception shall be raised. |

10348 These functions may fail if:

10349 **MX** **Range Error** The value of *x* is subnormal.

10350 If the integer expression (math\_errhandling & MATH\_ERRNO) is non-zero, |  
 10351 then *errno* shall be set to [ERANGE]. If the integer expression |  
 10352 (math\_errhandling & MATH\_ERREXCEPT) is non-zero, then the underflow |  
 10353 floating-point exception shall be raised. |

10354 **EXAMPLES**

10355 None.

10356 **APPLICATION USAGE**10357 The value of  $\text{expm1}(x)$  may be more accurate than  $\text{exp}(x)-1.0$  for small values of  $x$ .10358 The  $\text{expm1}()$  and  $\text{log1p}()$  functions are useful for financial calculations of  $((1+x)^n-1)/x$ , namely:10359  $\text{expm1}(n * \text{log1p}(x))/x$ 10360 when  $x$  is very small (for example, when calculating small daily interest rates). These functions  
10361 also simplify writing accurate inverse hyperbolic functions.10362 For IEEE Std 754-1985 **double**,  $709.8 < x$  implies  $\text{expm1}(x)$  has overflowed.10363 On error, the expressions (`math_errhandling & MATH_ERRNO`) and (`math_errhandling &`  
10364 `MATH_ERREXCEPT`) are independent of each other, but at least one of them must be non-zero.10365 **RATIONALE**

10366 None.

10367 **FUTURE DIRECTIONS**

10368 None.

10369 **SEE ALSO**10370  $\text{exp}()$ ,  $\text{feclearexcept}()$ ,  $\text{fetestexcept}()$ ,  $\text{ilogb}()$ ,  $\text{log1p}()$ , the Base Definitions volume of |  
10371 IEEE Std 1003.1-200x, Section 4.18, Treatment of Error Conditions for Mathematical Functions, |  
10372 `<math.h>`10373 **CHANGE HISTORY**

10374 First released in Issue 4, Version 2.

10375 **Issue 5**

10376 Moved from X/OPEN UNIX extension to BASE.

10377 **Issue 6**10378 The  $\text{expm1f}()$  and  $\text{expm1l}()$  functions are added for alignment with the ISO/IEC 9899:1999  
10379 standard.10380 The  $\text{expm1}()$  function is no longer marked as an extension. |10381 The DESCRIPTION, RETURN VALUE, ERRORS, and APPLICATION USAGE sections are |  
10382 revised to align with the ISO/IEC 9899:1999 standard.10383 IEC 60559:1989 standard floating-point extensions over the ISO/IEC 9899:1999 standard are  
10384 marked.

10385 **NAME**

10386 fabs, fabsf, fabsl — absolute value function

10387 **SYNOPSIS**

10388 #include <math.h>

10389 double fabs(double x);

10390 float fabsf(float x);

10391 long double fabsl(long double x);

10392 **DESCRIPTION**

10393 **CX** The functionality described on this reference page is aligned with the ISO C standard. Any  
10394 conflict between the requirements described here and the ISO C standard is unintentional. This  
10395 volume of IEEE Std 1003.1-200x defers to the ISO C standard.

10396 These functions shall compute the absolute value of their argument  $x$ ,  $|x|$ .

10397 **RETURN VALUE**

10398 Upon successful completion, these functions shall return the absolute value of  $x$ .

10399 **MX** If  $x$  is NaN, a NaN shall be returned.

10400 If  $x$  is  $\pm 0$ ,  $+0$  shall be returned.

10401 If  $x$  is  $\pm\text{Inf}$ ,  $+\text{Inf}$  shall be returned.

10402 **ERRORS**

10403 No errors are defined.

10404 **EXAMPLES**

10405 None.

10406 **APPLICATION USAGE**

10407 None.

10408 **RATIONALE**

10409 None.

10410 **FUTURE DIRECTIONS**

10411 None.

10412 **SEE ALSO**

10413 *isnan()*, the Base Definitions volume of IEEE Std 1003.1-200x, <math.h>

10414 **CHANGE HISTORY**

10415 First released in Issue 1. Derived from Issue 1 of the SVID.

10416 **Issue 5**

10417 The DESCRIPTION is updated to indicate how an application should check for an error. This  
10418 text was previously published in the APPLICATION USAGE section.

10419 **Issue 6**

10420 The *fabsf()* and *fabsl()* functions are added for alignment with the ISO/IEC 9899:1999 standard.

10421 The DESCRIPTION, RETURN VALUE, ERRORS, and APPLICATION USAGE sections are  
10422 revised to align with the ISO/IEC 9899:1999 standard.

10423 IEC 60559:1989 standard floating-point extensions over the ISO/IEC 9899:1999 standard are  
10424 marked.

10425 **NAME**

10426 fattach — attach a STREAMS-based file descriptor to a file in the file system name space  
 10427 (STREAMS)

10428 **SYNOPSIS**

```
10429 XSR #include <stropts.h>
```

```
10430 int fattach(int fildev, const char *path);
```

10431

10432 **DESCRIPTION**

10433 The *fattach()* function shall attach a STREAMS-based file descriptor to a file, effectively |  
 10434 associating a pathname with *fildev*. The application shall ensure that the *fildev* argument is a |  
 10435 valid open file descriptor associated with a STREAMS file. The *path* argument points to a |  
 10436 pathname of an existing file. The application shall have the appropriate privileges, or is the |  
 10437 owner of the file named by *path* and has write permission. A successful call to *fattach()* shall |  
 10438 cause all pathnames that name the file named by *path* to name the STREAMS file associated with |  
 10439 *fildev*, until the STREAMS file is detached from the file. A STREAMS file can be attached to more |  
 10440 than one file and can have several pathnames associated with it. |

10441 The attributes of the named STREAMS file shall be initialized as follows: the permissions, user |  
 10442 ID, group ID, and times are set to those of the file named by *path*, the number of links is set to 1, |  
 10443 and the size and device identifier are set to those of the STREAMS file associated with *fildev*. If |  
 10444 any attributes of the named STREAMS file are subsequently changed (for example, by *chmod()*), |  
 10445 neither the attributes of the underlying file nor the attributes of the STREAMS file to which *fildev* |  
 10446 refers shall be affected.

10447 File descriptors referring to the underlying file, opened prior to an *fattach()* call, shall continue to |  
 10448 refer to the underlying file.

10449 **RETURN VALUE**

10450 Upon successful completion, *fattach()* shall return 0. Otherwise, -1 shall be returned and *errno* |  
 10451 set to indicate the error.

10452 **ERRORS**

10453 The *fattach()* function shall fail if:

- |                         |                |   |
|-------------------------|----------------|---|
| 10454<br>10455<br>10456 | [EACCES]       | Search permission is denied for a component of the path prefix, or the process is the owner of <i>path</i> but does not have write permissions on the file named by <i>path</i> . |
| 10457                   | [EBADF]        | The <i>fildev</i> argument is not a valid open file descriptor.   |
| 10458<br>10459          | [EBUSY]        | The file named by <i>path</i> is currently a mount point or has a STREAMS file attached to it.  |
| 10460<br>10461          | [ELOOP]        | A loop exists in symbolic links encountered during resolution of the <i>path</i> argument.  |
| 10462<br>10463<br>10464 | [ENAMETOOLONG] | The size of <i>path</i> exceeds {PATH_MAX} or a component of <i>path</i> is longer than {NAME_MAX}.   |
| 10465                   | [ENOENT]       | A component of <i>path</i> does not name an existing file or <i>path</i> is an empty string.  |
| 10466                   | [ENOTDIR]      | A component of the path prefix is not a directory.  |
| 10467<br>10468          | [EPERM]        | The effective user ID of the process is not the owner of the file named by <i>path</i> and the process does not have appropriate privilege.                                       |

- 10469 The *fattach()* function may fail if:
- 10470 [EINVAL] The *fdes* argument does not refer to a STREAMS file.
  - 10471 [ELOOP] More than {SYMLOOP\_MAX} symbolic links were encountered during  
10472 resolution of the *path* argument.
  - 10473 [ENAMETOOLONG]  
10474 Pathname resolution of a symbolic link produced an intermediate result |  
10475 whose length exceeds {PATH\_MAX}.
  - 10476 [EXDEV] A link to a file on another file system was attempted.

10477 **EXAMPLES**10478 **Attaching a File Descriptor to a File**

10479 In the following example, *fd* refers to an open STREAMS file. The call to *fattach()* associates this  
10480 STREAM with the file */tmp/named-STREAM*, such that any future calls to open */tmp/named-*  
10481 *STREAM*, prior to breaking the attachment via a call to *fdetach()*, will instead create a new file  
10482 handle referring to the STREAMS file associated with *fd*.

```
10483 #include <stropts.h>
10484 ...
10485     int fd;
10486     char *filename = "/tmp/named-STREAM";
10487     int ret;
10488
10489     ret = fattach(fd, filename);
```

10489 **APPLICATION USAGE**

10490 The *fattach()* function behaves similarly to the traditional *mount()* function in the way a file is  
10491 temporarily replaced by the root directory of the mounted file system. In the case of *fattach()*, the  
10492 replaced file need not be a directory and the replacing file is a STREAMS file.

10493 **RATIONALE**

10494 The file attributes of a file which has been the subject of an *fattach()* call are specifically set |  
10495 because of an artefact of the original implementation. The internal mechanism was the same as |  
10496 for the *mount()* function. Since *mount()* is typically only applied to directories, the effects when |  
10497 applied to a regular file are a little surprising, especially as regards the link count which rigidly |  
10498 remains one, even if there were several links originally and despite the fact that all original links |  
10499 refer to the STREAM as long as the *fattach()* remains in effect. |

10500 **FUTURE DIRECTIONS**

10501 None.

10502 **SEE ALSO**

10503 *fdetach()*, *isastream()*, the Base Definitions volume of IEEE Std 1003.1-200x, <stropts.h>

10504 **CHANGE HISTORY**

10505 First released in Issue 4, Version 2.

10506 **Issue 5**

10507 Moved from X/OPEN UNIX extension to BASE.

10508 The [EXDEV] error is added to the list of optional errors in the ERRORS section.

10509 **Issue 6**

- 10510 This function is marked as part of the XSI STREAMS Option Group.
- 10511 The DESCRIPTION is updated to avoid use of the term “must” for application requirements.
- 10512 The wording of the mandatory [ELOOP] error condition is updated, and a second optional
- 10513 [ELOOP] error condition is added.



10514 **NAME**

10515 fchdir — change working directory

10516 **SYNOPSIS**

10517 XSI #include &lt;unistd.h&gt;

10518 int fchdir(int *fildev*);

10519

10520 **DESCRIPTION**10521 The *fchdir()* function shall be equivalent to *chdir()* except that the directory that is to be the new |  
10522 current working directory is specified by the file descriptor *fildev*.10523 A conforming application can obtain a file descriptor for a file of type directory using *open()*,  
10524 provided that the file status flags and access modes do not contain O\_WRONLY or O\_RDWR.10525 **RETURN VALUE**10526 Upon successful completion, *fchdir()* shall return 0. Otherwise, it shall return -1 and set *errno* to  
10527 indicate the error. On failure the current working directory shall remain unchanged.10528 **ERRORS**10529 The *fchdir()* function shall fail if:10530 [EACCES] Search permission is denied for the directory referenced by *fildev*.10531 [EBADF] The *fildev* argument is not an open file descriptor.10532 [ENOTDIR] The open file descriptor *fildev* does not refer to a directory.10533 The *fchdir()* may fail if:10534 [EINTR] A signal was caught during the execution of *fchdir()*.

10535 [EIO] An I/O error occurred while reading from or writing to the file system.

10536 **EXAMPLES**

10537 None.

10538 **APPLICATION USAGE**

10539 None.

10540 **RATIONALE**

10541 None.

10542 **FUTURE DIRECTIONS**

10543 None.

10544 **SEE ALSO**10545 *chdir()*, the Base Definitions volume of IEEE Std 1003.1-200x, <unistd.h>10546 **CHANGE HISTORY**

10547 First released in Issue 4, Version 2.

10548 **Issue 5**

10549 Moved from X/OPEN UNIX extension to BASE.

10550 **NAME**

10551 fchmod — change mode of a file

10552 **SYNOPSIS**

10553 #include &lt;sys/stat.h&gt;

10554 int fchmod(int *fildev*, mode\_t *mode*);10555 **DESCRIPTION**10556 The *fchmod()* function shall be equivalent to *chmod()* except that the file whose permissions are  
10557 changed is specified by the file descriptor *fildev*.10558 SHM If *fildev* references a shared memory object, the *fchmod()* function need only affect the S\_IRUSR,  
10559 S\_IWUSR, S\_IRGRP, S\_IWGRP, S\_IROTH, and S\_IWOTH file permission bits.10560 TYM If *fildev* references a typed memory object, the behavior of *fchmod()* is unspecified. |10561 If *fildev* refers to a socket, the behavior of *fchmod()* is unspecified. |10562 XSR If *fildev* refers to a STREAM (which is *fattach()*-ed into the file system name space) the call  
10563 returns successfully, doing nothing. |10564 **RETURN VALUE**10565 Upon successful completion, *fchmod()* shall return 0. Otherwise, it shall return -1 and set *errno* to  
10566 indicate the error.10567 **ERRORS**10568 The *fchmod()* function shall fail if:10569 [EBADF] The *fildev* argument is not an open file descriptor.10570 [EPERM] The effective user ID does not match the owner of the file and the process  
10571 does not have appropriate privilege.10572 [EROFS] The file referred to by *fildev* resides on a read-only file system.10573 The *fchmod()* function may fail if:10574 XSI [EINTR] The *fchmod()* function was interrupted by a signal.10575 XSI [EINVAL] The value of the *mode* argument is invalid.10576 [EINVAL] The *fildev* argument refers to a pipe and the implementation disallows  
10577 execution of *fchmod()* on a pipe.10578 **EXAMPLES**10579 **Changing the Current Permissions for a File**10580 The following example shows how to change the permissions for a file named */home/cnd/mod1*  
10581 so that the owner and group have read/write/execute permissions, but the world only has  
10582 read/write permissions.

10583 #include &lt;sys/stat.h&gt;

10584 #include &lt;fcntl.h&gt;

10585 mode\_t mode;

10586 int fildev;

10587 ...

10588 fildev = open("/home/cnd/mod1", O\_RDWR);

10589 fchmod(fildev, S\_IRWXU | S\_IRWXG | S\_IROTH | S\_IWOTH);

**10590 APPLICATION USAGE**

10591 None.

**10592 RATIONALE**

10593 None.

**10594 FUTURE DIRECTIONS**

10595 None.

**10596 SEE ALSO**

10597 *chmod()*, *chown()*, *creat()*, *fcntl()*, *fstatvfs()*, *mknod()*, *open()*, *read()*, *stat()*, *write()*, the Base  
10598 Definitions volume of IEEE Std 1003.1-200x, <sys/stat.h>

**10599 CHANGE HISTORY**

10600 First released in Issue 4, Version 2.

**10601 Issue 5**

10602 Moved from X/OPEN UNIX extension to BASE and aligned with *fchmod()* in the POSIX  
10603 Realtime Extension. Specifically, the second paragraph of the DESCRIPTION is added and a  
10604 second instance of [EINVAL] is defined in the list of optional errors.

**10605 Issue 6**

10606 The DESCRIPTION is updated for alignment with IEEE Std 1003.1j-2000 by stating that *fchmod()*  
10607 behavior is unspecified for typed memory objects.

## 10608 NAME

10609 fchown — change owner and group of a file

## 10610 SYNOPSIS

10611 #include &lt;unistd.h&gt;

10612 int fchown(int *fildev*, uid\_t *owner*, gid\_t *group*);

## 10613 DESCRIPTION

10614 The *fchown()* function shall be equivalent to *chown()* except that the file whose owner and group  
10615 are changed is specified by the file descriptor *fildev*.

## 10616 RETURN VALUE

10617 Upon successful completion, *fchown()* shall return 0. Otherwise, it shall return -1 and set *errno* to  
10618 indicate the error.

## 10619 ERRORS

10620 The *fchown()* function shall fail if:10621 [EBADF] The *fildev* argument is not an open file descriptor.10622 [EPERM] The effective user ID does not match the owner of the file or the process does  
10623 not have appropriate privilege and `_POSIX_CHOWN_RESTRICTED` indicates  
10624 that such privilege is required.10625 [EROFS] The file referred to by *fildev* resides on a read-only file system.10626 The *fchown()* function may fail if:10627 [EINVAL] The owner or group ID is not a value supported by the implementation. The  
10628 XSR *fildev* argument refers to a pipe or socket or an *fcntl()*-ed STREAM and the  
10629 implementation disallows execution of *fchown()* on a pipe.

10630 [EIO] A physical I/O error has occurred.

10631 [EINTR] The *fchown()* function was interrupted by a signal which was caught.

## 10632 EXAMPLES

10633 **Changing the Current Owner of a File**10634 The following example shows how to change the owner of a file named `/home/cnd/mod1` to  
10635 “jones” and the group to “cnd”.10636 The numeric value for the user ID is obtained by extracting the user ID from the user database  
10637 entry associated with “jones”. Similarly, the numeric value for the group ID is obtained by  
10638 extracting the group ID from the group database entry associated with “cnd”. This example  
10639 assumes the calling program has appropriate privileges.

10640 #include &lt;sys/types.h&gt;

10641 #include &lt;unistd.h&gt;

10642 #include &lt;fcntl.h&gt;

10643 #include &lt;pwd.h&gt;

10644 #include &lt;grp.h&gt;

10645 struct passwd \*pwd;

10646 struct group \*grp;

10647 int fildev;

10648 ...

10649 fildev = open("/home/cnd/mod1", O\_RDWR);

10650 pwd = getpwnam("jones");

```
10651     grp = getgrnam("cnd");
10652     fchown(fildes, pwd->pw_uid, grp->gr_gid);
```

**10653 APPLICATION USAGE**

10654 None.

**10655 RATIONALE**

10656 None.

**10657 FUTURE DIRECTIONS**

10658 None.

**10659 SEE ALSO**

10660 *chown()*, the Base Definitions volume of IEEE Std 1003.1-200x, <**unistd.h**>

**10661 CHANGE HISTORY**

10662 First released in Issue 4, Version 2.

**10663 Issue 5**

10664 Moved from X/OPEN UNIX extension to BASE. |

**10665 Issue 6**

10666 The following changes were made to align with the IEEE P1003.1a draft standard:

10667 • Clarification is added that a call to *fchown()* may not be allowed on a pipe. |

10668 The *fchown()* function is now defined as mandatory. |

10669 **NAME**

10670 `fclose` — close a stream

10671 **SYNOPSIS**

10672 `#include <stdio.h>`

10673 `int fclose(FILE *stream);`

10674 **DESCRIPTION**

10675 CX The functionality described on this reference page is aligned with the ISO C standard. Any  
 10676 conflict between the requirements described here and the ISO C standard is unintentional. This  
 10677 volume of IEEE Std 1003.1-200x defers to the ISO C standard.

10678 The `fclose()` function shall cause the stream pointed to by `stream` to be flushed and the associated  
 10679 file to be closed. Any unwritten buffered data for the stream shall be written to the file; any  
 10680 unread buffered data shall be discarded. Whether or not the call succeeds, the stream shall be  
 10681 disassociated from the file and any buffer set by the `setbuf()` or `setvbuf()` function shall be  
 10682 disassociated from the stream. If the associated buffer was automatically allocated, it shall be  
 10683 deallocated.

10684 CX The `fclose()` function shall mark for update the `st_ctime` and `st_mtime` fields of the underlying file,  
 10685 if the stream was writable, and if buffered data remains that has not yet been written to the file.  
 10686 The `fclose()` function shall perform the equivalent of a `close()` on the file descriptor that is  
 10687 associated with the stream pointed to by `stream`.

10688 After the call to `fclose()`, any use of `stream` results in undefined behavior.

10689 **RETURN VALUE**

10690 CX Upon successful completion, `fclose()` shall return 0; otherwise, it shall return EOF and set `errno` to  
 10691 indicate the error.

10692 **ERRORS**

10693 The `fclose()` function shall fail if:

10694 CX [EAGAIN] The O\_NONBLOCK flag is set for the file descriptor underlying `stream` and the  
 10695 process would be delayed in the write operation.

10696 CX [EBADF] The file descriptor underlying stream is not valid.

10697 CX [EFBIG] An attempt was made to write a file that exceeds the maximum file size.

10698 XSI [EFBIG] An attempt was made to write a file that exceeds the process' file size limit.

10699 CX [EFBIG] The file is a regular file and an attempt was made to write at or beyond the  
 10700 offset maximum associated with the corresponding stream.

10701 CX [EINTR] The `fclose()` function was interrupted by a signal.

10702 CX [EIO] The process is a member of a background process group attempting to write  
 10703 to its controlling terminal, TOSTOP is set, the process is neither ignoring nor  
 10704 blocking SIGTTOU, and the process group of the process is orphaned. This  
 10705 error may also be returned under implementation-defined conditions.

10706 CX [ENOSPC] There was no free space remaining on the device containing the file.

10707 CX [EPIPE] An attempt is made to write to a pipe or FIFO that is not open for reading by  
 10708 any process. A SIGPIPE signal shall also be sent to the thread.

10709 The `fclose()` function may fail if:

10710 CX [ENXIO] A request was made of a nonexistent device, or the request was outside the  
 10711 capabilities of the device.

10712 **EXAMPLES**

10713 None.

10714 **APPLICATION USAGE**

10715 None.

10716 **RATIONALE**

10717 None.

10718 **FUTURE DIRECTIONS**

10719 None.

10720 **SEE ALSO**

10721 *close()*, *fopen()*, *getrlimit()*, *ulimit()*, the Base Definitions volume of IEEE Std 1003.1-200x,  
10722 `<stdio.h>`

10723 **CHANGE HISTORY**

10724 First released in Issue 1. Derived from Issue 1 of the SVID.

10725 **Issue 5**

10726 Large File Summit extensions are added.

10727 **Issue 6**

10728 Extensions beyond the ISO C standard are now marked.

10729 The following new requirements on POSIX implementations derive from alignment with the  
10730 Single UNIX Specification:

- 10731 • The [EFBIG] error is added as part of the large file support extensions.
- 10732 • The [ENXIO] optional error condition is added.

10733 The DESCRIPTION is updated to note that the stream and any buffer are disassociated whether  
10734 or not the call succeeds. This is for alignment with the ISO/IEC 9899:1999 standard.

## 10735 NAME

10736 fcntl — file control

## 10737 SYNOPSIS

10738 OH #include &lt;unistd.h&gt;

10739 #include &lt;fcntl.h&gt;

10740 int fcntl(int *fildes*, int *cmd*, ...);

## 10741 DESCRIPTION

10742 The *fcntl()* function shall perform the operations described below on open files. The *fildes* |  
 10743 argument is a file descriptor.

10744 The available values for *cmd* are defined in <fcntl.h> and are as follows: |

10745 **F\_DUPFD** Return a new file descriptor which shall be the lowest numbered available |  
 10746 (that is, not already open) file descriptor greater than or equal to the third |  
 10747 argument, *arg*, taken as an integer of type **int**. The new file descriptor shall |  
 10748 refer to the same open file description as the original file descriptor, and shall |  
 10749 share any locks. The FD\_CLOEXEC flag associated with the new file |  
 10750 descriptor shall be cleared to keep the file open across calls to one of the *exec* |  
 10751 functions.

10752 **F\_GETFD** Get the file descriptor flags defined in <fcntl.h> that are associated with the |  
 10753 file descriptor *fildes*. File descriptor flags are associated with a single file |  
 10754 descriptor and do not affect other file descriptors that refer to the same file.

10755 **F\_SETFD** Set the file descriptor flags defined in <fcntl.h>, that are associated with *fildes*, |  
 10756 to the third argument, *arg*, taken as type **int**. If the FD\_CLOEXEC flag in the |  
 10757 third argument is 0, the file shall remain open across the *exec* functions; |  
 10758 otherwise, the file shall be closed upon successful execution of one of the *exec* |  
 10759 functions.

10760 **F\_GETFL** Get the file status flags and file access modes, defined in <fcntl.h>, for the file |  
 10761 description associated with *fildes*. The file access modes can be extracted from |  
 10762 the return value using the mask O\_ACCMODE, which is defined in <fcntl.h>. |  
 10763 File status flags and file access modes are associated with the file description |  
 10764 and do not affect other file descriptors that refer to the same file with different |  
 10765 open file descriptions.

10766 **F\_SETFL** Set the file status flags, defined in <fcntl.h>, for the file description associated |  
 10767 with *fildes* from the corresponding bits in the third argument, *arg*, taken as |  
 10768 type **int**. Bits corresponding to the file access mode and the file creation flags, |  
 10769 as defined in <fcntl.h>, that are set in *arg* shall be ignored. If any bits in *arg* |  
 10770 other than those mentioned here are changed by the application, the result is |  
 10771 unspecified.

10772 **F\_GETOWN** If *fildes* refers to a socket, get the process or process group ID specified to |  
 10773 receive SIGURG signals when out-of-band data is available. Positive values |  
 10774 indicate a process ID; negative values, other than -1, indicate a process group |  
 10775 ID. If *fildes* does not refer to a socket, the results are unspecified.

10776 **F\_SETOWN** If *fildes* refers to a socket, set the process or process group ID specified to |  
 10777 receive SIGURG signals when out-of-band data is available, using the value of |  
 10778 the third argument, *arg*, taken as type **int**. Positive values indicate a process |  
 10779 ID; negative values, other than -1, indicate a process group ID. If *fildes* does |  
 10780 not refer to a socket, the results are unspecified.



10781 The following values for *cmd* are available for advisory record locking. Record locking shall be |  
 10782 supported for regular files, and may be supported for other files. |

10783 **F\_GETLK** Get the first lock which blocks the lock description pointed to by the third |  
 10784 argument, *arg*, taken as a pointer to type **struct flock**, defined in `<fcntl.h>`. |  
 10785 The information retrieved shall overwrite the information passed to *fcntl()* in |  
 10786 the structure **flock**. If no lock is found that would prevent this lock from |  
 10787 being created, then the structure shall be left unchanged except for the lock |  
 10788 type which shall be set to **F\_UNLCK**.

10789 **F\_SETLK** Set or clear a file segment lock according to the lock description pointed to by |  
 10790 the third argument, *arg*, taken as a pointer to type **struct flock**, defined in |  
 10791 `<fcntl.h>`. **F\_SETLK** can establish shared (or read) locks (**F\_RDLCK**) or |  
 10792 exclusive (or write) locks (**F\_WRLCK**), as well as to remove either type of lock |  
 10793 (**F\_UNLCK**). **F\_RDLCK**, **F\_WRLCK**, and **F\_UNLCK** are defined in `<fcntl.h>`. |  
 10794 If a shared or exclusive lock cannot be set, *fcntl()* shall return immediately |  
 10795 with a return value of `-1`.

10796 **F\_SETLKW** This command shall be equivalent to **F\_SETLK** except that if a shared or |  
 10797 exclusive lock is blocked by other locks, the thread shall wait until the request |  
 10798 can be satisfied. If a signal that is to be caught is received while *fcntl()* is |  
 10799 waiting for a region, *fcntl()* shall be interrupted. Upon return from the signal |  
 10800 handler, *fcntl()* shall return `-1` with *errno* set to `[EINTR]`, and the lock |  
 10801 operation shall not be done.

10802 Additional implementation-defined values for *cmd* may be defined in `<fcntl.h>`. Their names |  
 10803 shall start with **F\_**.

10804 When a shared lock is set on a segment of a file, other processes shall be able to set shared locks |  
 10805 on that segment or a portion of it. A shared lock prevents any other process from setting an |  
 10806 exclusive lock on any portion of the protected area. A request for a shared lock shall fail if the |  
 10807 file descriptor was not opened with read access.

10808 An exclusive lock shall prevent any other process from setting a shared lock or an exclusive lock |  
 10809 on any portion of the protected area. A request for an exclusive lock shall fail if the file |  
 10810 descriptor was not opened with write access.

10811 The structure **flock** describes the type (*l\_type*), starting offset (*l\_whence*), relative offset (*l\_start*), |  
 10812 size (*l\_len*), and process ID (*l\_pid*) of the segment of the file to be affected.

10813 The value of *l\_whence* is **SEEK\_SET**, **SEEK\_CUR**, or **SEEK\_END**, to indicate that the relative |  
 10814 offset *l\_start* bytes shall be measured from the start of the file, current position, or end of the file, |  
 10815 respectively. The value of *l\_len* is the number of consecutive bytes to be locked. The value of |  
 10816 *l\_len* may be negative (where the definition of **off\_t** permits negative values of *l\_len*). The *l\_pid* |  
 10817 field is only used with **F\_GETLK** to return the process ID of the process holding a blocking lock. |  
 10818 After a successful **F\_GETLK** request, when a blocking lock is found, the values returned in the |  
 10819 **flock** structure shall be as follows:

10820 *l\_type* Type of blocking lock found.

10821 *l\_whence* **SEEK\_SET**.

10822 *l\_start* Start of the blocking lock.

10823 *l\_len* Length of the blocking lock.

10824 *l\_pid* Process ID of the process that holds the blocking lock.

10825 If the command is F\_SETLKW and the process must wait for another process to release a lock,  
 10826 then the range of bytes to be locked shall be determined before the *fcntl()* function blocks. If the  
 10827 file size or file descriptor seek offset change while *fcntl()* is blocked, this shall not affect the  
 10828 range of bytes locked.

10829 If *l\_len* is positive, the area affected shall start at *l\_start* and end at *l\_start+l\_len-1*. If *l\_len* is |  
 10830 negative, the area affected shall start at *l\_start+l\_len* and end at *l\_start-1*. Locks may start and |  
 10831 extend beyond the current end of a file, but shall not extend before the beginning of the file. A |  
 10832 lock shall be set to extend to the largest possible value of the file offset for that file by setting |  
 10833 *l\_len* to 0. If such a lock also has *l\_start* set to 0 and *l\_whence* is set to SEEK\_SET, the whole file  
 10834 shall be locked.

10835 There shall be at most one type of lock set for each byte in the file. Before a successful return  
 10836 from an F\_SETLK or an F\_SETLKW request when the calling process has previously existing  
 10837 locks on bytes in the region specified by the request, the previous lock type for each byte in the  
 10838 specified region shall be replaced by the new lock type. As specified above under the  
 10839 descriptions of shared locks and exclusive locks, an F\_SETLK or an F\_SETLKW request  
 10840 (respectively) shall fail or block when another process has existing locks on bytes in the specified  
 10841 region and the type of any of those locks conflicts with the type specified in the request.

10842 All locks associated with a file for a given process shall be removed when a file descriptor for  
 10843 that file is closed by that process or the process holding that file descriptor terminates. Locks are  
 10844 not inherited by a child process.

10845 A potential for deadlock occurs if a process controlling a locked region is put to sleep by  
 10846 attempting to lock another process' locked region. If the system detects that sleeping until a  
 10847 locked region is unlocked would cause a deadlock, *fcntl()* shall fail with an [EDEADLK] error.

10848 An unlock (F\_UNLCK) request in which *l\_len* is non-zero and the offset of the last byte of the |  
 10849 requested segment is the maximum value for an object of type **off\_t**, when the process has an |  
 10850 existing lock in which *l\_len* is 0 and which includes the last byte of the requested segment, shall |  
 10851 be treated as a request to unlock from the start of the requested segment with an *l\_len* equal to 0. |  
 10852 Otherwise, an unlock (F\_UNLCK) request shall attempt to unlock only the requested segment. |

10853 SHM When the file descriptor *fdes* refers to a shared memory object, the behavior of *fcntl()* shall be |  
 10854 the same as for a regular file except the effect of the following values for the argument *cmd* shall |  
 10855 be unspecified: F\_SETFL, F\_GETLK, F\_SETLK, and F\_SETLKW.

10856 TYM If *fdes* refers to a typed memory object, the result of the *fcntl()* function is unspecified. |

#### 10857 RETURN VALUE

10858 Upon successful completion, the value returned shall depend on *cmd* as follows:

10859	F_DUPFD	A new file descriptor.
10860	F_GETFD	Value of flags defined in <fcntl.h>. The return value shall not be negative.
10861	F_SETFD	Value other than -1.
10862	F_GETFL	Value of file status flags and access modes. The return value is not negative.
10863	F_SETFL	Value other than -1.
10864	F_GETLK	Value other than -1.
10865	F_SETLK	Value other than -1.
10866	F_SETLKW	Value other than -1.
10867	F_GETOWN	Value of the socket owner process or process group; this will not be -1.

- 10868           F\_SETOWN        Value other than `-1`.
- 10869           Otherwise, `-1` shall be returned and *errno* set to indicate the error.
- 10870 **ERRORS**
- 10871           The *fcntl()* function shall fail if:
- 10872           [EACCES] or [EAGAIN]
- 10873                        The *cmd* argument is `F_SETLK`; the type of lock (*l\_type*) is a shared (`F_RDLCK`) or exclusive (`F_WRLCK`) lock and the segment of a file to be locked is already exclusive-locked by another process, or the type is an exclusive lock and some portion of the segment of a file to be locked is already shared-locked or exclusive-locked by another process.
- 10874
- 10875
- 10876
- 10877
- 10878           [EBADF]        The *fildev* argument is not a valid open file descriptor, or the argument *cmd* is `F_SETLK` or `F_SETLKW`, the type of lock, *l\_type*, is a shared lock (`F_RDLCK`), and *fildev* is not a valid file descriptor open for reading, or the type of lock *l\_type*, is an exclusive lock (`F_WRLCK`), and *fildev* is not a valid file descriptor open for writing.
- 10879
- 10880
- 10881
- 10882
- 10883           [EINTR]        The *cmd* argument is `F_SETLKW` and the function was interrupted by a signal.
- 10884           [EINVAL]        The *cmd* argument is invalid, or the *cmd* argument is `F_DUPFD` and *arg* is negative or greater than or equal to `{OPEN_MAX}`, or the *cmd* argument is `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.
- 10885
- 10886
- 10887
- 10888           [EMFILE]        The argument *cmd* is `F_DUPFD` and `{OPEN_MAX}` file descriptors are currently open in the calling process, or no file descriptors greater than or equal to *arg* are available.
- 10889
- 10890
- 10891           [ENOLCK]        The argument *cmd* is `F_SETLK` or `F_SETLKW` and satisfying the lock or unlock request would result in the number of locked regions in the system exceeding a system-imposed limit.
- 10892
- 10893
- 10894           [E\_OVERFLOW]    One of the values to be returned cannot be represented correctly.
- 10895           [E\_OVERFLOW]    The *cmd* argument is `F_GETLK`, `F_SETLK`, or `F_SETLKW` and the smallest or, 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`.
- 10896
- 10897
- 10898           The *fcntl()* function may fail if:
- 10899           [EDEADLK]        The *cmd* argument is `F_SETLKW`, the lock is blocked by some lock from another process and putting the calling process to sleep, waiting for that lock to become free would cause a deadlock.
- 10900
- 10901

10902 **EXAMPLES**

10903           None.

10904 **APPLICATION USAGE**

10905           None.

10906 **RATIONALE**

10907           The ellipsis in the SYNOPSIS is the syntax specified by the ISO C standard for a variable number of arguments. It is used because System V uses pointers for the implementation of file locking functions.

10908

10909

10910           The *arg* values to `F_GETFD`, `F_SETFD`, `F_GETFL`, and `F_SETFL` all represent flag values to allow for future growth. Applications using these functions should do a read-modify-write operation

10911

10912 on them, rather than assuming that only the values defined by this volume of  
10913 IEEE Std 1003.1-200x are valid. It is a common error to forget this, particularly in the case of  
10914 F\_SETFD.

10915 This volume of IEEE Std 1003.1-200x permits concurrent read and write access to file data using  
10916 the *fcntl()* function; this is a change from the 1984 /usr/group standard and early proposals.  
10917 Without concurrency controls, this feature may not be fully utilized without occasional loss of  
10918 data.

10919 Data losses occur in several ways. One case occurs when several processes try to update the  
10920 same record, without sequencing controls; several updates may occur in parallel and the last  
10921 writer “wins”. Another case is a bit-tree or other internal list-based database that is undergoing  
10922 reorganization. Without exclusive use to the tree segment by the updating process, other reading  
10923 processes chance getting lost in the database when the index blocks are split, condensed,  
10924 inserted, or deleted. While *fcntl()* is useful for many applications, it is not intended to be overly  
10925 general and does not handle the bit-tree example well.

10926 This facility is only required for regular files because it is not appropriate for many devices such  
10927 as terminals and network connections.

10928 Since *fcntl()* works with “any file descriptor associated with that file, however it is obtained”,  
10929 the file descriptor may have been inherited through a *fork()* or *exec* operation and thus may  
10930 affect a file that another process also has open.

10931 The use of the open file description to identify what to lock requires extra calls and presents  
10932 problems if several processes are sharing an open file description, but there are too many  
10933 implementations of the existing mechanism for this volume of IEEE Std 1003.1-200x to use  
10934 different specifications.

10935 Another consequence of this model is that closing any file descriptor for a given file (whether or  
10936 not it is the same open file description that created the lock) causes the locks on that file to be  
10937 relinquished for that process. Equivalently, any close for any file/process pair relinquishes the  
10938 locks owned on that file for that process. But note that while an open file description may be  
10939 shared through *fork()*, locks are not inherited through *fork()*. Yet locks may be inherited through  
10940 one of the *exec* functions.

10941 The identification of a machine in a network environment is outside of the scope of this volume  
10942 of IEEE Std 1003.1-200x. Thus, an *L\_sysid* member, such as found in System V, is not included in  
10943 the locking structure.

10944 Changing of lock types can result in a previously locked region being split into smaller regions. |

10945 Mandatory locking was a major feature of the 1984 /usr/group standard.

10946 For advisory file record locking to be effective, all processes that have access to a file must  
10947 cooperate and use the advisory mechanism before doing I/O on the file. Enforcement-mode  
10948 record locking is important when it cannot be assumed that all processes are cooperating. For  
10949 example, if one user uses an editor to update a file at the same time that a second user executes  
10950 another process that updates the same file and if only one of the two processes is using advisory  
10951 locking, the processes are not cooperating. Enforcement-mode record locking would protect  
10952 against accidental collisions.

10953 Secondly, advisory record locking requires a process using locking to bracket each I/O operation  
10954 with lock (or test) and unlock operations. With enforcement-mode file and record locking, a  
10955 process can lock the file once and unlock when all I/O operations have been completed.  
10956 Enforcement-mode record locking provides a base that can be enhanced; for example, with  
10957 sharable locks. That is, the mechanism could be enhanced to allow a process to lock a file so  
10958 other processes could read it, but none of them could write it.

- 10959 Mandatory locks were omitted for several reasons:
- 10960 1. Mandatory lock setting was done by multiplexing the set-group-ID bit in most  
10961 implementations; this was confusing, at best.
  - 10962 2. The relationship to file truncation as supported in 4.2 BSD was not well specified.
  - 10963 3. Any publicly readable file could be locked by anyone. Many historical implementations  
10964 keep the password database in a publicly readable file. A malicious user could thus  
10965 prohibit logins. Another possibility would be to hold open a long-distance telephone line.
  - 10966 4. Some demand-paged historical implementations offer memory mapped files, and  
10967 enforcement cannot be done on that type of file.
- 10968 Since sleeping on a region is interrupted with any signal, *alarm()* may be used to provide a  
10969 timeout facility in applications requiring it. This is useful in deadlock detection. Since |  
10970 implementation of full deadlock detection is not always feasible, the [EDEADLK] error was |  
10971 made optional.
- 10972 **FUTURE DIRECTIONS**
- 10973 None.
- 10974 **SEE ALSO**
- 10975 *close()*, *exec*, *open()*, *sigaction()*, the Base Definitions volume of IEEE Std 1003.1-200x, <**fcntl.h**>,  
10976 <**signal.h**>, <**unistd.h**>
- 10977 **CHANGE HISTORY**
- 10978 First released in Issue 1. Derived from Issue 1 of the SVID.
- 10979 **Issue 5**
- 10980 The DESCRIPTION is updated for alignment with the POSIX Realtime Extension and the POSIX  
10981 Threads Extension.
- 10982 Large File Summit extensions are added.
- 10983 **Issue 6**
- 10984 In the SYNOPSIS, the optional include of the <**sys/types.h**> header is removed.
- 10985 The following new requirements on POSIX implementations derive from alignment with the  
10986 Single UNIX Specification:
- 10987 • The requirement to include <**sys/types.h**> has been removed. Although <**sys/types.h**> was  
10988 required for conforming implementations of previous POSIX specifications, it was not  
10989 required for UNIX applications.
  - 10990 • In the DESCRIPTION, sentences describing behavior when *L\_len* is negative are now  
10991 mandated, and the description of unlock (F\_UNLOCK) when *L\_len* is non-negative is  
10992 mandated.
  - 10993 • In the ERRORS section, the [EINVAL] error condition has the case mandated when the *cmd* is  
10994 invalid, and two [EOVERFLOW] error conditions are added.
- 10995 The F\_GETOWN and F\_SETOWN values are added for sockets.
- 10996 The following changes were made to align with the IEEE P1003.1a draft standard:
- 10997 • Clarification is added that the extent of the bytes locked is determined prior to the blocking  
10998 action.
- 10999 The DESCRIPTION is updated for alignment with IEEE Std 1003.1j-2000 by specifying that  
11000 *fcntl()* results are unspecified for typed memory objects.

11001

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

11002 **NAME**11003       fcvt — convert a floating-point number to a string (**LEGACY**)11004 **SYNOPSIS**

11005 xSI     #include &lt;stdlib.h&gt;

11006       char \*fcvt(double *value*, int *ndigit*, int \*restrict *decpt*,  
11007           int \*restrict *sign*);

11008

11009 **DESCRIPTION**11010       Refer to *ecvt()*.

11011 **NAME**11012 fdatasync — synchronize the data of a file (**REALTIME**)11013 **SYNOPSIS**11014 SIO `#include <unistd.h>`11015 `int fdatasync(int fildes);`

11016

11017 **DESCRIPTION**11018 The *fdatasync()* function shall force all currently queued I/O operations associated with the file  
11019 indicated by file descriptor *fil*des to the synchronized I/O completion state.11020 The functionality shall be equivalent to *fsync()* with the symbol `_POSIX_SYNCHRONIZED_IO` |  
11021 defined, with the exception that all I/O operations shall be completed as defined for |  
11022 synchronized I/O data integrity completion.11023 **RETURN VALUE**11024 If successful, the *fdatasync()* function shall return the value 0; otherwise, the function shall return  
11025 the value `-1` and set *errno* to indicate the error. If the *fdatasync()* function fails, outstanding I/O  
11026 operations are not guaranteed to have been completed.11027 **ERRORS**11028 The *fdatasync()* function shall fail if:11029 [EBADF] The *fil*des argument is not a valid file descriptor open for writing.

11030 [EINVAL] This implementation does not support synchronized I/O for this file.

11031 In the event that any of the queued I/O operations fail, *fdatasync()* shall return the error  
11032 conditions defined for *read()* and *write()*.11033 **EXAMPLES**

11034 None.

11035 **APPLICATION USAGE**

11036 None.

11037 **RATIONALE**

11038 None.

11039 **FUTURE DIRECTIONS**

11040 None.

11041 **SEE ALSO**11042 *ai*o\_11043 *fsync()*, *fcntl()*, *fsync()*, *open()*, *read()*, *write()*, the Base Definitions volume of  
IEEE Std 1003.1-200x, `<unistd.h>`11044 **CHANGE HISTORY**

11045 First released in Issue 5. Included for alignment with the POSIX Realtime Extension.

11046 **Issue 6**11047 The [ENOSYS] error condition has been removed as stubs need not be provided if an  
11048 implementation does not support the Synchronized Input and Output option.11049 The *fdatasync()* function is marked as part of the Synchronized Input and Output option.



11050 **NAME**11051 fdetach — detach a name from a STREAMS-based file descriptor (**STREAMS**)11052 **SYNOPSIS**

11053 XSR #include &lt;stropts.h&gt;

11054 int fdetach(const char \*path);

11055

11056 **DESCRIPTION**

11057 The *fdetach()* function shall detach a STREAMS-based file from the file to which it was attached |  
 11058 by a previous call to *fattach()*. The *path* argument points to the pathname of the attached |  
 11059 STREAMS file. The process shall have appropriate privileges or be the owner of the file. A |  
 11060 successful call to *fdetach()* shall cause all pathnames that named the attached STREAMS file to |  
 11061 again name the file to which the STREAMS file was attached. All subsequent operations on *path* |  
 11062 shall operate on the underlying file and not on the STREAMS file.

11063 All open file descriptions established while the STREAMS file was attached to the file referenced |  
 11064 by *path* shall still refer to the STREAMS file after the *fdetach()* has taken effect.

11065 If there are no open file descriptors or other references to the STREAMS file, then a successful |  
 11066 call to *fdetach()* shall have be equivalent to performing the last *close()* on the attached file. |

11067 **RETURN VALUE**

11068 Upon successful completion, *fdetach()* shall return 0; otherwise, it shall return -1 and set *errno* to |  
 11069 indicate the error.

11070 **ERRORS**11071 The *fdetach()* function shall fail if:

11072 [EACCES] Search permission is denied on a component of the path prefix.

11073 [EINVAL] The *path* argument names a file that is not currently attached.

11074 [ELOOP] A loop exists in symbolic links encountered during resolution of the *path* |  
 11075 argument.

11076 [ENAMETOOLONG]

11077 The size of a pathname exceeds {PATH\_MAX} or a pathname component is |  
 11078 longer than {NAME\_MAX}. |

11079 [ENOENT] A component of *path* does not name an existing file or *path* is an empty string.

11080 [ENOTDIR] A component of the path prefix is not a directory.

11081 [EPERM] The effective user ID is not the owner of *path* and the process does not have |  
 11082 appropriate privileges.

11083 The *fdetach()* function may fail if:

11084 [ELOOP] More than {SYMLOOP\_MAX} symbolic links were encountered during |  
 11085 resolution of the *path* argument.

11086 [ENAMETOOLONG]

11087 Pathname resolution of a symbolic link produced an intermediate result |  
 11088 whose length exceeds {PATH\_MAX}.

11089 **EXAMPLES**11090 **Detaching a File**

11091 The following example detaches the STREAMS-based file **/tmp/named-STREAM** from the file to  
11092 which it was attached by a previous, successful call to *fattach()*. Subsequent calls to open this  
11093 file refer to the underlying file, not to the STREAMS file.

```
11094 #include <stropts.h>
11095 ...
11096     char *filename = "/tmp/named-STREAM";
11097     int ret;
11098     ret = fdetach(filename);
```

11099 **APPLICATION USAGE**

11100 None.

11101 **RATIONALE**

11102 None.

11103 **FUTURE DIRECTIONS**

11104 None.

11105 **SEE ALSO**

11106 *fattach()*, the Base Definitions volume of IEEE Std 1003.1-200x, <**stropts.h**>

11107 **CHANGE HISTORY**

11108 First released in Issue 4, Version 2.

11109 **Issue 5**

11110 Moved from X/OPEN UNIX extension to BASE.

11111 **Issue 6**

11112 The DESCRIPTION is updated to avoid use of the term “must” for application requirements.

11113 The wording of the mandatory [ELOOP] error condition is updated, and a second optional  
11114 [ELOOP] error condition is added.

11115 **NAME**

11116 `fdim`, `fdimf`, `fdiml` — compute positive difference between two floating-point numbers

11117 **SYNOPSIS**

11118 `#include <math.h>`

11119 `double fdim(double x, double y);`

11120 `float fdimf(float x, float y);`

11121 `long double fdiml(long double x, long double y);`

11122 **DESCRIPTION**

11123 **CX** The functionality described on this reference page is aligned with the ISO C standard. Any  
 11124 conflict between the requirements described here and the ISO C standard is unintentional. This  
 11125 volume of IEEE Std 1003.1-200x defers to the ISO C standard.

11126 These functions shall determine the positive difference between their arguments. If  $x$  is greater  
 11127 than  $y$ ,  $x-y$  is returned. If  $x$  is less than or equal to  $y$ ,  $+0$  is returned.

11128 An application wishing to check for error situations should set *errno* to zero and call  
 11129 *feclearexcept*(*FE\_ALL\_EXCEPT*) before calling these functions. On return, if *errno* is non-zero or  
 11130 *fetestexcept*(*FE\_INVALID* | *FE\_DIVBYZERO* | *FE\_OVERFLOW* | *FE\_UNDERFLOW*) is non-  
 11131 zero, an error has occurred.

11132 **RETURN VALUE**

11133 Upon successful completion, these functions shall return the positive difference value.

11134 If  $x-y$  is positive and overflows, a range error shall occur and *fdim*(), *fdimf*(), and *fdiml*() shall  
 11135 return the value of the macro *HUGE\_VAL*, *HUGE\_VALF*, and *HUGE\_VALL*, respectively.

11136 **XSI** If  $x-y$  is positive and underflows, a range error may occur, and either  $(x-y)$  (if representable), or  
 11137  $0.0$  (if supported), or an implementation-defined value shall be returned.

11138 **MX** If  $x$  or  $y$  is NaN, a NaN shall be returned.

11139 **ERRORS**

11140 The *fdim*() function shall fail if:

11141 **Range Error** The result overflows.

11142 If the integer expression (*math\_errhandling* & *MATH\_ERRNO*) is non-zero, |  
 11143 then *errno* shall be set to [ERANGE]. If the integer expression |  
 11144 (*math\_errhandling* & *MATH\_ERREXCEPT*) is non-zero, then the overflow |  
 11145 floating-point exception shall be raised. |

11146 The *fdim*() function may fail if:

11147 **Range Error** The result underflows.

11148 If the integer expression (*math\_errhandling* & *MATH\_ERRNO*) is non-zero, |  
 11149 then *errno* shall be set to [ERANGE]. If the integer expression |  
 11150 (*math\_errhandling* & *MATH\_ERREXCEPT*) is non-zero, then the underflow |  
 11151 floating-point exception shall be raised. |

11152 **EXAMPLES**

11153       None.

11154 **APPLICATION USAGE**

11155       On implementations supporting IEEE Std 754-1985,  $x-y$  cannot underflow, and hence the 0.0  
11156       return value is shaded as an extension for implementations supporting the XSI extension rather  
11157       than an MX extension.

11158       On error, the expressions `(math_errhandling & MATH_ERRNO)` and `(math_errhandling &`  
11159       `MATH_ERREXCEPT)` are independent of each other, but at least one of them must be non-zero.

11160 **RATIONALE**

11161       None.

11162 **FUTURE DIRECTIONS**

11163       None.

11164 **SEE ALSO**

11165       *feclearexcept()*, *fetestexcept()*, *fmax()*, *fmin()*, the Base Definitions volume of IEEE Std 1003.1-200x, |  
11166       Section 4.18, Treatment of Error Conditions for Mathematical Functions, <math.h> |

11167 **CHANGE HISTORY**

11168       First released in Issue 6. Derived from the ISO/IEC 9899:1999 standard.

11169 **NAME**

11170 fdopen — associate a stream with a file descriptor

11171 **SYNOPSIS**

11172 cx #include &lt;stdio.h&gt;

11173 FILE \*fdopen(int *fil-des*, const char \**mode*);

11174

11175 **DESCRIPTION**11176 The *fdopen()* function shall associate a stream with a file descriptor.11177 The *mode* argument is a character string having one of the following values:11178 *r* or *rb* Open a file for reading.11179 *w* or *wb* Open a file for writing.11180 *a* or *ab* Open a file for writing at end of file.11181 *r+* or *rb+* or *r+b* Open a file for update (reading and writing).11182 *w+* or *wb+* or *w+b* Open a file for update (reading and writing).11183 *a+* or *ab+* or *a+b* Open a file for update (reading and writing) at end of file.11184 The meaning of these flags is exactly as specified in *fopen()*, except that modes beginning with *w* |  
11185 shall not cause truncation of the file. |11186 Additional values for the *mode* argument may be supported by an implementation.11187 The application shall ensure that the mode of the stream as expressed by the *mode* argument is  
11188 allowed by the file access mode of the open file description to which *fil-des* refers. The file  
11189 position indicator associated with the new stream is set to the position indicated by the file  
11190 offset associated with the file descriptor.11191 The error and end-of-file indicators for the stream shall be cleared. The *fdopen()* function may  
11192 cause the *st\_atime* field of the underlying file to be marked for update.11193 SHM If *fil-des* refers to a shared memory object, the result of the *fdopen()* function is unspecified.11194 TYM If *fil-des* refers to a typed memory object, the result of the *fdopen()* function is unspecified.11195 The *fdopen()* function shall preserve the offset maximum previously set for the open file  
11196 description corresponding to *fil-des*.11197 **RETURN VALUE**11198 Upon successful completion, *fdopen()* shall return a pointer to a stream; otherwise, a null pointer  
11199 shall be returned and *errno* set to indicate the error.11200 **ERRORS**11201 The *fdopen()* function may fail if:11202 [EBADF] The *fil-des* argument is not a valid file descriptor.11203 [EINVAL] The *mode* argument is not a valid mode.

11204 [EMFILE] {FOPEN\_MAX} streams are currently open in the calling process.

11205 [EMFILE] {STREAM\_MAX} streams are currently open in the calling process.

11206 [ENOMEM] Insufficient space to allocate a buffer.

11207 **EXAMPLES**

11208 None.

11209 **APPLICATION USAGE**

11210 File descriptors are obtained from calls like *open()*, *dup()*, *creat()*, or *pipe()*, which open files but  
 11211 do not return streams.

11212 **RATIONALE**

11213 The file descriptor may have been obtained from *open()*, *creat()*, *pipe()*, *dup()*, or *fcntl()*;  
 11214 inherited through *fork()* or *exec*; or perhaps obtained by implementation-defined means, such as  
 11215 the 4.3 BSD *socket()* call.

11216 The meanings of the *mode* arguments of *fdopen()* and *fopen()* differ. With *fdopen()*, open for write  
 11217 (*w* or *w+*) does not truncate, and append (*a* or *a+*) cannot create for writing. The *mode* argument  
 11218 formats that include *a b* are allowed for consistency with the ISO C standard function *fopen()*.  
 11219 The *b* has no effect on the resulting stream. Although not explicitly required by this volume of  
 11220 IEEE Std 1003.1-200x, a good implementation of append (*a*) mode would cause the *O\_APPEND*  
 11221 flag to be set.

11222 **FUTURE DIRECTIONS**

11223 None.

11224 **SEE ALSO**

11225 *fclose()*, *fopen()*, *open()*, the Base Definitions volume of IEEE Std 1003.1-200x, <*stdio.h*>, Section  
 11226 2.5.1 (on page 485)

11227 **CHANGE HISTORY**

11228 First released in Issue 1. Derived from Issue 1 of the SVID.

11229 **Issue 5**

11230 The DESCRIPTION is updated for alignment with the POSIX Realtime Extension.

11231 Large File Summit extensions are added.

11232 **Issue 6**

11233 The following new requirements on POSIX implementations derive from alignment with the  
 11234 Single UNIX Specification:

- 11235 • In the DESCRIPTION, the use and setting of the *mode* argument are changed to include  
 11236 binary streams.
- 11237 • In the DESCRIPTION, text is added for large file support to indicate setting of the offset  
 11238 maximum in the open file description.
- 11239 • All errors identified in the ERRORS section are added.
- 11240 • In the DESCRIPTION, text is added that the *fdopen()* function may cause *st\_atime* to be  
 11241 updated.

11242 The following changes were made to align with the IEEE P1003.1a draft standard:

- 11243 • Clarification is added that it is the responsibility of the application to ensure that the mode is  
 11244 compatible with the open file descriptor.

11245 The DESCRIPTION is updated for alignment with IEEE Std 1003.1j-2000 by specifying that  
 11246 *fdopen()* results are unspecified for typed memory objects.

11247 **NAME**

11248           feclearexcept — clear floating-point exception

11249 **SYNOPSIS**

11250           #include &lt;fenv.h&gt;

11251           int feclearexcept(int *excepts*);11252 **DESCRIPTION**

11253 *cx*       The functionality described on this reference page is aligned with the ISO C standard. Any  
11254       conflict between the requirements described here and the ISO C standard is unintentional. This  
11255       volume of IEEE Std 1003.1-200x defers to the ISO C standard.

11256       The *feclearexcept()* function shall attempt to clear the supported floating-point exceptions |  
11257       represented by *excepts*. |

11258 **RETURN VALUE**

11259       If the argument is zero or if all the specified exceptions were successfully cleared, *feclearexcept()* |  
11260       shall return zero. Otherwise, it shall return a non-zero value. |

11261 **ERRORS**

11262       No errors are defined.

11263 **EXAMPLES**

11264       None.

11265 **APPLICATION USAGE**

11266       None.

11267 **RATIONALE**

11268       None.

11269 **FUTURE DIRECTIONS**

11270       None.

11271 **SEE ALSO**

11272       *fegetexceptflag()*, *feraiseexcept()*, *fesetexceptflag()*, *fetestexcept()*, the Base Definitions volume of  
11273       IEEE Std 1003.1-200x, <**fenv.h**>

11274 **CHANGE HISTORY**

11275       First released in Issue 6. Derived from the ISO/IEC 9899:1999 standard. |

11276       ISO/IEC 9899:1999 standard, Technical Corrigendum No. 1 is incorporated. |

11277 **NAME**

11278 fegetenv, fesetenv — get and set current floating-point environment

11279 **SYNOPSIS**

11280 #include &lt;fenv.h&gt;

11281 int fegetenv(fenv\_t \*envp);

11282 int 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 attempt to store the current floating-point environment in the object  
11288 pointed to by *envp*.

11289 The *fesetenv()* function shall attempt to establish the floating-point environment represented by  
11290 the object pointed to by *envp*. The argument *envp* shall point to an object set by a call to  
11291 *fegetenv()* or *feholdexcept()*, or equal a floating-point environment macro. The *fesetenv()* function  
11292 does not raise floating-point exceptions, but only installs the state of the floating-point status  
11293 flags represented through its argument.

11294 **RETURN VALUE**

11295 If the representation was successfully stored, *fegetenv()* shall return zero. Otherwise, it shall  
11296 return a non-zero value. If the environment was successfully established, *fesetenv()* shall return  
11297 zero. Otherwise, it shall return a non-zero value.

11298 **ERRORS**

11299 No errors are defined.

11300 **EXAMPLES**

11301 None.

11302 **APPLICATION USAGE**

11303 None.

11304 **RATIONALE**

11305 None.

11306 **FUTURE DIRECTIONS**

11307 None.

11308 **SEE ALSO**11309 *feholdexcept()*, *feupdateenv()*, the Base Definitions volume of IEEE Std 1003.1-200x, <fenv.h>11310 **CHANGE HISTORY**

11311 First released in Issue 6. Derived from the ISO/IEC 9899:1999 standard.

11312 ISO/IEC 9899:1999 standard, Technical Corrigendum No. 1 is incorporated.



11313 **NAME**

11314 fegetexceptflag, fesetexceptflag — get and set floating-point status flags

11315 **SYNOPSIS**

11316 #include &lt;fenv.h&gt;

11317 int fegetexceptflag(fexcept\_t \*flagp, int excepts);

11318 int fesetexceptflag(const fexcept\_t \*flagp, int excepts);

11319 **DESCRIPTION**

11320 **CX** The functionality described on this reference page is aligned with the ISO C standard. Any  
 11321 conflict between the requirements described here and the ISO C standard is unintentional. This  
 11322 volume of IEEE Std 1003.1-200x defers to the ISO C standard.

11323 The *fegetexceptflag()* function shall attempt to store an implementation-defined representation of  
 11324 the states of the floating-point status flags indicated by the argument *excepts* in the object  
 11325 pointed to by the argument *flagp*.

11326 The *fesetexceptflag()* function shall attempt to set the floating-point status flags indicated by the  
 11327 argument *excepts* to the states stored in the object pointed to by *flagp*. The value pointed to by  
 11328 *flagp* shall have been set by a previous call to *fegetexceptflag()* whose second argument  
 11329 represented at least those floating-point exceptions represented by the argument *excepts*. This  
 11330 function does not raise floating-point exceptions, but only sets the state of the flags.

11331 **RETURN VALUE**

11332 If the representation was successfully stored, *fegetexceptflag()* shall return zero. Otherwise, it  
 11333 shall return a non-zero value. If the *excepts* argument is zero or if all the specified exceptions  
 11334 were successfully set, *fesetexceptflag()* shall return zero. Otherwise, it shall return a non-zero  
 11335 value.

11336 **ERRORS**

11337 No errors are defined.

11338 **EXAMPLES**

11339 None.

11340 **APPLICATION USAGE**

11341 None.

11342 **RATIONALE**

11343 None.

11344 **FUTURE DIRECTIONS**

11345 None.

11346 **SEE ALSO**

11347 *feclearexcept()*, *feraiseexcept()*, *fetestexcept()*, the Base Definitions volume of IEEE Std 1003.1-200x,  
 11348 <fenv.h>

11349 **CHANGE HISTORY**

11350 First released in Issue 6. Derived from the ISO/IEC 9899:1999 standard.

11351 ISO/IEC 9899:1999 standard, Technical Corrigendum No. 1 is incorporated.

11352 **NAME**

11353 fegetround, fesetround — get and set current rounding direction

11354 **SYNOPSIS**

```
11355     #include <fenv.h>
11356     int fegetround(void);
11357     int fesetround(int round);
```

11358 **DESCRIPTION**

11359 **CX** The functionality described on this reference page is aligned with the ISO C standard. Any  
 11360 conflict between the requirements described here and the ISO C standard is unintentional. This  
 11361 volume of IEEE Std 1003.1-200x defers to the ISO C standard.

11362 The *fegetround()* function shall get the current rounding direction.

11363 The *fesetround()* function shall establish the rounding direction represented by its argument  
 11364 *round*. If the argument is not equal to the value of a rounding direction macro, the rounding  
 11365 direction is not changed.

11366 **RETURN VALUE**

11367 The *fegetround()* function shall return the value of the rounding direction macro representing the  
 11368 current rounding direction or a negative value if there is no such rounding direction macro or  
 11369 the current rounding direction is not determinable.

11370 The *fesetround()* function shall return a zero value if and only if the requested rounding direction  
 11371 was established.

11372 **ERRORS**

11373 No errors are defined.

11374 **EXAMPLES**

11375 The following example saves, sets, and restores the rounding direction, reporting an error and  
 11376 aborting if setting the rounding direction fails:

```
11377     #include <fenv.h>
11378     #include <assert.h>
11379     void f(int round_dir)
11380     {
11381         #pragma STDC FENV_ACCESS ON
11382         int save_round;
11383         int setround_ok;
11384         save_round = fegetround();
11385         setround_ok = fesetround(round_dir);
11386         assert(setround_ok == 0);
11387         /* ... */
11388         fesetround(save_round);
11389         /* ... */
11390     }
```

11391 **APPLICATION USAGE**

11392 None.

11393 **RATIONALE**

11394 None.

11395 **FUTURE DIRECTIONS**

11396           None.

11397 **SEE ALSO**

11398           The Base Definitions volume of IEEE Std 1003.1-200x, <fenv.h>

11399 **CHANGE HISTORY**

11400           First released in Issue 6. Derived from the ISO/IEC 9899: 1999 standard. |

11401           ISO/IEC 9899: 1999 standard, Technical Corrigendum No. 1 is incorporated. |

11402 **NAME**

11403 feholdexcept — save current floating-point environment

11404 **SYNOPSIS**

11405 #include <fenv.h>

11406 int feholdexcept(fenv\_t \*envp);

11407 **DESCRIPTION**

11408 **CX** The functionality described on this reference page is aligned with the ISO C standard. Any  
11409 conflict between the requirements described here and the ISO C standard is unintentional. This  
11410 volume of IEEE Std 1003.1-200x defers to the ISO C standard.

11411 The *feholdexcept()* function shall save the current floating-point environment in the object  
11412 pointed to by *envp*, clear the floating-point status flags, and then install a non-stop (continue on  
11413 floating-point exceptions) mode, if available, for all floating-point exceptions.

11414 **RETURN VALUE**

11415 The *feholdexcept()* function shall return zero if and only if non-stop floating-point exception  
11416 handling was successfully installed.

11417 **ERRORS**

11418 No errors are defined.

11419 **EXAMPLES**

11420 None.

11421 **APPLICATION USAGE**

11422 None.

11423 **RATIONALE**

11424 The *feholdexcept()* function should be effective on typical IEC 60559:1989 standard  
11425 implementations which have the default non-stop mode and at least one other mode for trap  
11426 handling or aborting. If the implementation provides only the non-stop mode, then installing the  
11427 non-stop mode is trivial.

11428 **FUTURE DIRECTIONS**

11429 None.

11430 **SEE ALSO**

11431 *fegetenv()*, *fesetenv()*, *feupdateenv()*, the Base Definitions volume of IEEE Std 1003.1-200x,  
11432 <fenv.h>

11433 **CHANGE HISTORY**

11434 First released in Issue 6. Derived from the ISO/IEC 9899:1999 standard.

11435 **NAME**

11436 feof — test end-of-file indicator on a stream

11437 **SYNOPSIS**

11438 #include &lt;stdio.h&gt;

11439 int feof(FILE \**stream*);11440 **DESCRIPTION**

11441 cx The functionality described on this reference page is aligned with the ISO C standard. Any  
11442 conflict between the requirements described here and the ISO C standard is unintentional. This  
11443 volume of IEEE Std 1003.1-200x defers to the ISO C standard.

11444 The *feof()* function shall test the end-of-file indicator for the stream pointed to by *stream*.11445 **RETURN VALUE**11446 The *feof()* function shall return non-zero if and only if the end-of-file indicator is set for *stream*.11447 **ERRORS**

11448 No errors are defined.

11449 **EXAMPLES**

11450 None.

11451 **APPLICATION USAGE**

11452 None.

11453 **RATIONALE**

11454 None.

11455 **FUTURE DIRECTIONS**

11456 None.

11457 **SEE ALSO**11458 *clearerr()*, *ferror()*, *fopen()*, the Base Definitions volume of IEEE Std 1003.1-200x, <stdio.h>11459 **CHANGE HISTORY**

11460 First released in Issue 1. Derived from Issue 1 of the SVID.

11461 **NAME**11462        `feraiseexcept` — raise floating-point exception11463 **SYNOPSIS**11464        `#include <fenv.h>`11465        `int feraiseexcept(int excepts);`11466 **DESCRIPTION**11467 **cx**        The functionality described on this reference page is aligned with the ISO C standard. Any  
11468        conflict between the requirements described here and the ISO C standard is unintentional. This  
11469        volume of IEEE Std 1003.1-200x defers to the ISO C standard.11470        The `feraiseexcept()` function shall attempt to raise the supported floating-point exceptions |  
11471        represented by the argument `excepts`. The order in which these floating-point exceptions are |  
11472        raised is unspecified. Whether the `feraiseexcept()` function additionally raises the inexact |  
11473        floating-point exception whenever it raises the overflow or underflow floating-point exception is |  
11474        implementation-defined.11475 **RETURN VALUE**11476        If the argument is zero or if all the specified exceptions were successfully raised, `feraiseexcept()` |  
11477        shall return zero. Otherwise, it shall return a non-zero value. |11478 **ERRORS**

11479        No errors are defined.

11480 **EXAMPLES**

11481        None.

11482 **APPLICATION USAGE**11483        The effect is intended to be similar to that of floating-point exceptions raised by arithmetic  
11484        operations. Hence, enabled traps for floating-point exceptions raised by this function are taken.11485 **RATIONALE**11486        Raising overflow or underflow is allowed to also raise inexact because on some architectures the  
11487        only practical way to raise an exception is to execute an instruction that has the exception as a  
11488        side effect. The function is not restricted to accept only valid coincident expressions for atomic  
11489        operations, so the function can be used to raise exceptions accrued over several operations.11490 **FUTURE DIRECTIONS**

11491        None.

11492 **SEE ALSO**11493        `feclearexcept()`, `fegetexceptflag()`, `fesetexceptflag()`, `fetestexcept()`, the Base Definitions volume of  
11494        IEEE Std 1003.1-200x, `<fenv.h>`11495 **CHANGE HISTORY**

11496        First released in Issue 6. Derived from the ISO/IEC 9899:1999 standard. |

11497        ISO/IEC 9899:1999 standard, Technical Corrigendum No. 1 is incorporated. |

11498 **NAME**

11499           ferror — test error indicator on a stream

11500 **SYNOPSIS**

11501           #include &lt;stdio.h&gt;

11502           int ferror(FILE \**stream*);11503 **DESCRIPTION**

11504 **CX**       The functionality described on this reference page is aligned with the ISO C standard. Any  
11505           conflict between the requirements described here and the ISO C standard is unintentional. This  
11506           volume of IEEE Std 1003.1-200x defers to the ISO C standard.

11507           The *ferror()* function shall test the error indicator for the stream pointed to by *stream*.11508 **RETURN VALUE**11509           The *ferror()* function shall return non-zero if and only if the error indicator is set for *stream*.11510 **ERRORS**

11511           No errors are defined.

11512 **EXAMPLES**

11513           None.

11514 **APPLICATION USAGE**

11515           None.

11516 **RATIONALE**

11517           None.

11518 **FUTURE DIRECTIONS**

11519           None.

11520 **SEE ALSO**11521           *clearerr()*, *feof()*, *fopen()*, the Base Definitions volume of IEEE Std 1003.1-200x, <stdio.h>11522 **CHANGE HISTORY**

11523           First released in Issue 1. Derived from Issue 1 of the SVID.

11524 **NAME**

11525       fesetenv — set current floating-point environment

11526 **SYNOPSIS**

11527       #include <fenv.h>

11528       int fesetenv(const fenv\_t \*envp);

11529 **DESCRIPTION**

11530       Refer to *fegetenv()*.



11531 **NAME**

11532       fesetexceptflag — set floating-point status flags

11533 **SYNOPSIS**

11534       #include &lt;fenv.h&gt;

11535       int fesetexceptflag(const fexcept\_t \*flagp, int excepts);

11536 **DESCRIPTION**11537       Refer to *fegetexceptflag()*.

11538 **NAME**

11539       fesetround — set current rounding direction

11540 **SYNOPSIS**

11541       #include <fenv.h>

11542       int fesetround(int *round*);

11543 **DESCRIPTION**

11544       Refer to *fegetround()*.

11545 **NAME**

11546       fetetestexcept — test floating-point exception flags

11547 **SYNOPSIS**

11548       #include &lt;fenv.h&gt;

11549       int fetetestexcept(int *excepts*);11550 **DESCRIPTION**

11551 cx       The functionality described on this reference page is aligned with the ISO C standard. Any  
 11552       conflict between the requirements described here and the ISO C standard is unintentional. This  
 11553       volume of IEEE Std 1003.1-200x defers to the ISO C standard.

11554       The *fetetestexcept()* function shall determine which of a specified subset of the floating-point  
 11555       exception flags are currently set. The *excepts* argument specifies the floating-point status flags to  
 11556       be queried.

11557 **RETURN VALUE**

11558       The *fetetestexcept()* function shall return the value of the bitwise-inclusive OR of the floating-point  
 11559       exception macros corresponding to the currently set floating-point exceptions included in  
 11560       *excepts*.

11561 **ERRORS**

11562       No errors are defined.

11563 **EXAMPLES**

11564       The following example calls function *f()* if an invalid exception is set, and then function *g()* if an  
 11565       overflow exception is set:

```

11566       #include <fenv.h>
11567       /* ... */
11568       {
11569           #pragma STDC FENV_ACCESS ON
11570           int set_excepts;
11571           feclearexcept(FE_INVALID | FE_OVERFLOW);
11572           // maybe raise exceptions
11573           set_excepts = fetetestexcept(FE_INVALID | FE_OVERFLOW);
11574           if (set_excepts & FE_INVALID) f();
11575           if (set_excepts & FE_OVERFLOW) g();
11576           /* ... */
11577       }
```

11578 **APPLICATION USAGE**

11579       None.

11580 **RATIONALE**

11581       None.

11582 **FUTURE DIRECTIONS**

11583       None.

11584 **SEE ALSO**

11585       *feclearexcept()*, *fegetexceptflag()*, *feraiseexcept()*, the Base Definitions volume of  
 11586       IEEE Std 1003.1-200x, <fenv.h>

11587 **CHANGE HISTORY**

11588 First released in Issue 6. Derived from the ISO/IEC 9899: 1999 standard.

11589 **NAME**

11590 feupdateenv — update floating-point environment

11591 **SYNOPSIS**

11592 #include &lt;fenv.h&gt;

11593 int feupdateenv(const fenv\_t \*envp);

11594 **DESCRIPTION**

11595 cx The functionality described on this reference page is aligned with the ISO C standard. Any  
 11596 conflict between the requirements described here and the ISO C standard is unintentional. This  
 11597 volume of IEEE Std 1003.1-200x defers to the ISO C standard.

11598 The *feupdateenv()* function shall attempt to save the currently raised floating-point exceptions in  
 11599 its automatic storage, attempt to install the floating-point environment represented by the object  
 11600 pointed to by *envp*, and then attempt to raise the saved floating-point exceptions. The argument  
 11601 *envp* shall point to an object set by a call to *feholdexcept()* or *fegetenv()*, or equal a floating-point  
 11602 environment macro.

11603 **RETURN VALUE**

11604 The *feupdateenv()* function shall return a zero value if and only if all the required actions were  
 11605 successfully carried out.

11606 **ERRORS**

11607 No errors are defined.

11608 **EXAMPLES**

11609 The following example shows sample code to hide spurious underflow floating-point  
 11610 exceptions:

```

11611 #include <fenv.h>
11612 double f(double x)
11613 {
11614     #pragma STDC FENV_ACCESS ON
11615     double result;
11616     fenv_t save_env;
11617     feholdexcept(&save_env);
11618     // compute result
11619     if (/* test spurious underflow */)
11620         feclearexcept(FE_UNDERFLOW);
11621     feupdateenv(&save_env);
11622     return result;
11623 }
```

11624 **APPLICATION USAGE**

11625 None.

11626 **RATIONALE**

11627 None.

11628 **FUTURE DIRECTIONS**

11629 None.

11630 **SEE ALSO**11631 *fegetenv()*, *feholdexcept()*, the Base Definitions volume of IEEE Std 1003.1-200x, <fenv.h>

11632 **CHANGE HISTORY**

- 11633           First released in Issue 6. Derived from the ISO/IEC 9899: 1999 standard. |
- 11634           ISO/IEC 9899: 1999 standard, Technical Corrigendum No. 1 is incorporated. |

11635 **NAME**

11636 fflush — flush a stream

11637 **SYNOPSIS**

11638 #include &lt;stdio.h&gt;

11639 int fflush(FILE \*stream);

11640 **DESCRIPTION**

11641 CX The functionality described on this reference page is aligned with the ISO C standard. Any  
 11642 conflict between the requirements described here and the ISO C standard is unintentional. This  
 11643 volume of IEEE Std 1003.1-200x defers to the ISO C standard.

11644 If *stream* points to an output stream or an update stream in which the most recent operation was  
 11645 CX not input, *fflush()* shall cause any unwritten data for that stream to be written to the file, and the  
 11646 *st\_ctime* and *st\_mtime* fields of the underlying file shall be marked for update.

11647 If *stream* is a null pointer, *fflush()* shall perform this flushing action on all streams for which the  
 11648 behavior is defined above.

11649 **RETURN VALUE**

11650 Upon successful completion, *fflush()* shall return 0; otherwise, it shall set the error indicator for  
 11651 CX the stream, return EOF, and set *errno* to indicate the error.

11652 **ERRORS**11653 The *fflush()* function shall fail if:

11654 CX [EAGAIN] The O\_NONBLOCK flag is set for the file descriptor underlying *stream* and the  
 11655 process would be delayed in the write operation.

11656 CX [EBADF] The file descriptor underlying *stream* is not valid.

11657 CX [EFBIG] An attempt was made to write a file that exceeds the maximum file size.

11658 XSI [EFBIG] An attempt was made to write a file that exceeds the process' file size limit.

11659 CX [EFBIG] The file is a regular file and an attempt was made to write at or beyond the  
 11660 offset maximum associated with the corresponding stream.

11661 CX [EINTR] The *fflush()* function was interrupted by a signal.

11662 CX [EIO] The process is a member of a background process group attempting to write  
 11663 to its controlling terminal, TOSTOP is set, the process is neither ignoring nor  
 11664 blocking SIGTTOU, and the process group of the process is orphaned. This  
 11665 error may also be returned under implementation-defined conditions.

11666 CX [ENOSPC] There was no free space remaining on the device containing the file.

11667 CX [EPIPE] An attempt is made to write to a pipe or FIFO that is not open for reading by  
 11668 any process. A SIGPIPE signal shall also be sent to the thread.

11669 The *fflush()* function may fail if:

11670 CX [ENXIO] A request was made of a nonexistent device, or the request was outside the  
 11671 capabilities of the device.

11672 **EXAMPLES**11673 **Sending Prompts to Standard Output**

11674 The following example uses *printf()* calls to print a series of prompts for information the user  
11675 must enter from standard input. The *fflush()* calls force the output to standard output. The  
11676 *fflush()* function is used because standard output is usually buffered and the prompt may not  
11677 immediately be printed on the output or terminal. The *gets()* calls read strings from standard  
11678 input and place the results in variables, for use later in the program

```
11679 #include <stdio.h>
11680 ...
11681 char user[100];
11682 char oldpasswd[100];
11683 char newpasswd[100];
11684 ...
11685 printf("User name: ");
11686 fflush(stdout);
11687 gets(user);

11688 printf("Old password: ");
11689 fflush(stdout);
11690 gets(oldpasswd);

11691 printf("New password: ");
11692 fflush(stdout);
11693 gets(newpasswd);
11694 ...
```

11695 **APPLICATION USAGE**

11696 None.

11697 **RATIONALE**

11698 Data buffered by the system may make determining the validity of the position of the current  
11699 file descriptor impractical. Thus, enforcing the repositioning of the file descriptor after *fflush()*  
11700 on streams open for *read()* is not mandated by IEEE Std 1003.1-200x.

11701 **FUTURE DIRECTIONS**

11702 None.

11703 **SEE ALSO**

11704 *getrlimit()*, *ulimit()*, the Base Definitions volume of IEEE Std 1003.1-200x, **<stdio.h>**

11705 **CHANGE HISTORY**

11706 First released in Issue 1. Derived from Issue 1 of the SVID.

11707 **Issue 5**

11708 Large File Summit extensions are added.

11709 **Issue 6**

11710 Extensions beyond the ISO C standard are now marked.

11711 The following new requirements on POSIX implementations derive from alignment with the  
11712 Single UNIX Specification:

- 11713 • The [EFBIG] error is added as part of the large file support extensions.
- 11714 • The [ENXIO] optional error condition is added.



11715        The RETURN VALUE section is updated to note that the error indicator shall be set for the  
11716        stream. This is for alignment with the ISO/IEC 9899:1999 standard.

11717 **NAME**

11718       ffs — find first set bit

11719 **SYNOPSIS**

11720 xSI       #include &lt;strings.h&gt;

11721       int ffs(int i);

11722

11723 **DESCRIPTION**11724       The *ffs()* function shall find the first bit set (beginning with the least significant bit) in *i*, and  
11725       return the index of that bit. Bits are numbered starting at one (the least significant bit).11726 **RETURN VALUE**11727       The *ffs()* function shall return the index of the first bit set. If *i* is 0, then *ffs()* shall return 0.11728 **ERRORS**

11729       No errors are defined.

11730 **EXAMPLES**

11731       None.

11732 **APPLICATION USAGE**

11733       None.

11734 **RATIONALE**

11735       None.

11736 **FUTURE DIRECTIONS**

11737       None.

11738 **SEE ALSO**

11739       The Base Definitions volume of IEEE Std 1003.1-200x, &lt;strings.h&gt;

11740 **CHANGE HISTORY**

11741       First released in Issue 4, Version 2.

11742 **Issue 5**

11743       Moved from X/OPEN UNIX extension to BASE.

## 11744 NAME

11745 fgetc — get a byte from a stream

## 11746 SYNOPSIS

11747 #include &lt;stdio.h&gt;

11748 int fgetc(FILE \*stream);

## 11749 DESCRIPTION

11750 CX The functionality described on this reference page is aligned with the ISO C standard. Any  
 11751 conflict between the requirements described here and the ISO C standard is unintentional. This  
 11752 volume of IEEE Std 1003.1-200x defers to the ISO C standard.

11753 If the end-of-file indicator for the input stream pointed to by *stream* is not set and a next byte is  
 11754 present, the *fgetc()* function shall obtain the next byte as an **unsigned char** converted to an **int**,  
 11755 from the input stream pointed to by *stream*, and advance the associated file position indicator for  
 11756 the stream (if defined). Since *fgetc()* operates on bytes, reading a character consisting of multiple  
 11757 bytes (or “a multi-byte character”) may require multiple calls to *fgetc()*.

11758 CX The *fgetc()* function may mark the *st\_atime* field of the file associated with *stream* for update. The  
 11759 *st\_atime* field shall be marked for update by the first successful execution of *fgetc()*, *fgets()*,  
 11760 *fgetwc()*, *fgetws()*, *fread()*, *fscanf()*, *getc()*, *getchar()*, *gets()*, or *scanf()* using *stream* that returns  
 11761 data not supplied by a prior call to *ungetc()* or *ungetwc()*.

## 11762 RETURN VALUE

11763 Upon successful completion, *fgetc()* shall return the next byte from the input stream pointed to  
 11764 by *stream*. If the end-of-file indicator for the stream is set, or if the stream is at end-of-file, the  
 11765 end-of-file indicator for the stream shall be set and *fgetc()* shall return EOF. If a read error occurs,  
 11766 CX the error indicator for the stream shall be set, *fgetc()* shall return EOF, and shall set *errno* to  
 11767 indicate the error.

## 11768 ERRORS

11769 The *fgetc()* function shall fail if data needs to be read and:

11770 CX [EAGAIN] The O\_NONBLOCK flag is set for the file descriptor underlying *stream* and the  
 11771 process would be delayed in the *fgetc()* operation.

11772 CX [EBADF] The file descriptor underlying *stream* is not a valid file descriptor open for  
 11773 reading.

11774 CX [EINTR] The read operation was terminated due to the receipt of a signal, and no data  
 11775 was transferred.

11776 CX [EIO] A physical I/O error has occurred, or the process is in a background process  
 11777 group attempting to read from its controlling terminal, and either the process  
 11778 is ignoring or blocking the SIGTTIN signal or the process group is orphaned.  
 11779 This error may also be generated for implementation-defined reasons.

11780 CX [EOVERFLOW] The file is a regular file and an attempt was made to read at or beyond the  
 11781 offset maximum associated with the corresponding stream.

11782 The *fgetc()* function may fail if:

11783 CX [ENOMEM] Insufficient storage space is available.

11784 CX [ENXIO] A request was made of a nonexistent device, or the request was outside the  
 11785 capabilities of the device.

11786 **EXAMPLES**

11787       None.

11788 **APPLICATION USAGE**

11789       If the integer value returned by *fgetc()* is stored into a variable of type **char** and then compared  
11790       against the integer constant EOF, the comparison may never succeed, because sign-extension of  
11791       a variable of type **char** on widening to integer is implementation-defined.

11792       The *ferror()* or *feof()* functions must be used to distinguish between an error condition and an  
11793       end-of-file condition.

11794 **RATIONALE**

11795       None.

11796 **FUTURE DIRECTIONS**

11797       None.

11798 **SEE ALSO**

11799       *feof()*, *ferror()*, *fopen()*, *getchar()*, *getc()*, the Base Definitions volume of IEEE Std 1003.1-200x,  
11800       <stdio.h>

11801 **CHANGE HISTORY**

11802       First released in Issue 1. Derived from Issue 1 of the SVID.

11803 **Issue 5**

11804       Large File Summit extensions are added.

11805 **Issue 6**

11806       Extensions beyond the ISO C standard are now marked.

11807       The following new requirements on POSIX implementations derive from alignment with the  
11808       Single UNIX Specification:

- 11809       • The [EIO] and [EOVERFLOW] mandatory error conditions are added.
- 11810       • The [ENOMEM] and [ENXIO] optional error conditions are added.

11811       The following changes are made for alignment with the ISO/IEC 9899:1999 standard:

- 11812       • The DESCRIPTION is updated to clarify the behavior when the end-of-file indicator for the  
11813       input stream is not set.
- 11814       • The RETURN VALUE section is updated to note that the error indicator shall be set for the  
11815       stream.

11816 **NAME**

11817           fgetpos — get current file position information

11818 **SYNOPSIS**

11819           #include &lt;stdio.h&gt;

11820           int fgetpos(FILE \*restrict stream, fpos\_t \*restrict pos);

11821 **DESCRIPTION**

11822 CX       The functionality described on this reference page is aligned with the ISO C standard. Any  
 11823       conflict between the requirements described here and the ISO C standard is unintentional. This  
 11824       volume of IEEE Std 1003.1-200x defers to the ISO C standard.

11825       The *fgetpos()* function shall store the current values of the parse state (if any) and file position  
 11826       indicator for the stream pointed to by *stream* in the object pointed to by *pos*. The value stored  
 11827       contains unspecified information usable by *fsetpos()* for repositioning the stream to its position  
 11828       at the time of the call to *fgetpos()*.

11829 **RETURN VALUE**

11830       Upon successful completion, *fgetpos()* shall return 0; otherwise, it shall return a non-zero value  
 11831       and set *errno* to indicate the error.

11832 **ERRORS**11833       The *fgetpos()* function shall fail if:

11834 CX       [E\_OVERFLOW]   The current value of the file position cannot be represented correctly in an  
 11835       object of type **fpos\_t**.

11836       The *fgetpos()* function may fail if:

11837 CX       [EBADF]        The file descriptor underlying *stream* is not valid.

11838 CX       [ESPIPE]       The file descriptor underlying *stream* is associated with a pipe, FIFO, or socket.  
 11839

11840 **EXAMPLES**

11841       None.

11842 **APPLICATION USAGE**

11843       None.

11844 **RATIONALE**

11845       None.

11846 **FUTURE DIRECTIONS**

11847       None.

11848 **SEE ALSO**11849       *fopen()*, *ftell()*, *rewind()*, *ungetc()*, the Base Definitions volume of IEEE Std 1003.1-200x, <stdio.h>11850 **CHANGE HISTORY**

11851       First released in Issue 4. Derived from the ISO C standard.

11852 **Issue 5**

11853       Large File Summit extensions are added.

11854 **Issue 6**

11855       Extensions beyond the ISO C standard are now marked.

11856       The following new requirements on POSIX implementations derive from alignment with the  
 11857       Single UNIX Specification:

- 11858       • The [EIO] mandatory error condition is added.
- 11859       • The [EBADF] and [ESPIPE] optional error conditions are added.
- 11860       An additional [ESPIPE] error condition is added for sockets.
- 11861       The prototype for *fgetpos()* is changed for alignment with the ISO/IEC 9899:1999 standard.

11862 **NAME**

11863       fgets — get a string from a stream

11864 **SYNOPSIS**

11865       #include &lt;stdio.h&gt;

11866       char \*fgets(char \*restrict *s*, int *n*, FILE \*restrict *stream*);11867 **DESCRIPTION**

11868 cx     The functionality described on this reference page is aligned with the ISO C standard. Any  
 11869     conflict between the requirements described here and the ISO C standard is unintentional. This  
 11870     volume of IEEE Std 1003.1-200x defers to the ISO C standard.

11871     The *fgets()* function shall read bytes from *stream* into the array pointed to by *s*, until *n*−1 bytes  
 11872     are read, or a <newline> is read and transferred to *s*, or an end-of-file condition is encountered.  
 11873     The string is then terminated with a null byte.

11874 cx     The *fgets()* function may mark the *st\_atime* field of the file associated with *stream* for update. The  
 11875     *st\_atime* field shall be marked for update by the first successful execution of *fgetc()*, *fgets()*,  
 11876     *fgetwc()*, *fgetws()*, *fread()*, *fscanf()*, *getc()*, *getchar()*, *gets()*, or *scanf()* using *stream* that returns  
 11877     data not supplied by a prior call to *ungetc()* or *ungetwc()*.

11878 **RETURN VALUE**

11879     Upon successful completion, *fgets()* shall return *s*. If the stream is at end-of-file, the end-of-file  
 11880     indicator for the stream shall be set and *fgets()* shall return a null pointer. If a read error occurs,  
 11881 cx     the error indicator for the stream shall be set, *fgets()* shall return a null pointer, and shall set  
 11882     *errno* to indicate the error.

11883 **ERRORS**11884       Refer to *fgetc()*.11885 **EXAMPLES**11886       **Reading Input**

11887     The following example uses *fgets()* to read each line of input. {LINE\_MAX}, which defines the  
 11888     maximum size of the input line, is defined in the <limits.h> header.

```
11889     #include <stdio.h>
11890     ...
11891     char line[LINE_MAX];
11892     ...
11893     while (fgets(line, LINE_MAX, fp) != NULL) {
11894         ...
11895     }
11896     ...
```

11897 **APPLICATION USAGE**

11898       None.

11899 **RATIONALE**

11900       None.

11901 **FUTURE DIRECTIONS**

11902       None.

11903 **SEE ALSO**

11904 *fopen()*, *fread()*, *gets()*, the Base Definitions volume of IEEE Std 1003.1-200x, <stdio.h>

11905 **CHANGE HISTORY**

11906 First released in Issue 1. Derived from Issue 1 of the SVID.

11907 **Issue 6**

11908 Extensions beyond the ISO C standard are now marked.

11909 The prototype for *fgets()* is changed for alignment with the ISO/IEC 9899:1999 standard.



11910 **NAME**

11911 fgetwc — get a wide-character code from a stream

11912 **SYNOPSIS**

11913 #include &lt;stdio.h&gt;

11914 #include &lt;wchar.h&gt;

11915 wint\_t fgetwc(FILE \*stream);

11916 **DESCRIPTION**

11917 CX The functionality described on this reference page is aligned with the ISO C standard. Any  
 11918 conflict between the requirements described here and the ISO C standard is unintentional. This  
 11919 volume of IEEE Std 1003.1-200x defers to the ISO C standard.

11920 The *fgetwc()* function shall obtain the next character (if present) from the input stream pointed to  
 11921 by *stream*, convert that to the corresponding wide-character code, and advance the associated  
 11922 file position indicator for the stream (if defined).

11923 If an error occurs, the resulting value of the file position indicator for the stream is unspecified. |

11924 CX The *fgetwc()* function may mark the *st\_atime* field of the file associated with *stream* for update.  
 11925 The *st\_atime* field shall be marked for update by the first successful execution of *fgetc()*, *fgets()*,  
 11926 *fgetwc()*, *fgetws()*, *fread()*, *fscanf()*, *getc()*, *getchar()*, *gets()*, or *scanf()* using *stream* that returns  
 11927 data not supplied by a prior call to *ungetc()* or *ungetwc()*.

11928 **RETURN VALUE**

11929 Upon successful completion, the *fgetwc()* function shall return the wide-character code of the  
 11930 character read from the input stream pointed to by *stream* converted to a type **wint\_t**. If the  
 11931 stream is at end-of-file, the end-of-file indicator for the stream shall be set and *fgetwc()* shall  
 11932 return WEOF. If a read error occurs, the error indicator for the stream shall be set, *fgetwc()* shall  
 11933 CX return WEOF, and shall set *errno* to indicate the error. If an encoding error occurs, the error |  
 11934 indicator for the stream shall be set, *fgetwc()* shall return WEOF, and shall set *errno* to indicate |  
 11935 the error. |

11936 **ERRORS**11937 The *fgetwc()* function shall fail if data needs to be read and:

11938 CX [EAGAIN] The O\_NONBLOCK flag is set for the file descriptor underlying *stream* and the  
 11939 process would be delayed in the *fgetwc()* operation.

11940 CX [EBADF] The file descriptor underlying *stream* is not a valid file descriptor open for  
 11941 reading. |

11942 [EILSEQ] The data obtained from the input stream does not form a valid character. |

11943 CX [EINTR] The read operation was terminated due to the receipt of a signal, and no data  
 11944 was transferred.

11945 CX [EIO] A physical I/O error has occurred, or the process is in a background process  
 11946 group attempting to read from its controlling terminal, and either the process  
 11947 is ignoring or blocking the SIGTTIN signal or the process group is orphaned.  
 11948 This error may also be generated for implementation-defined reasons.

11949 CX [EOVERFLOW] The file is a regular file and an attempt was made to read at or beyond the  
 11950 offset maximum associated with the corresponding stream.

11951 The *fgetwc()* function may fail if:

11952 CX [ENOMEM] Insufficient storage space is available.

11953 CX [ENXIO] A request was made of a nonexistent device, or the request was outside the  
11954 capabilities of the device.

#### 11955 EXAMPLES

11956 None.

#### 11957 APPLICATION USAGE

11958 The *ferror()* or *feof()* functions must be used to distinguish between an error condition and an  
11959 end-of-file condition.

#### 11960 RATIONALE

11961 None.

#### 11962 FUTURE DIRECTIONS

11963 None.

#### 11964 SEE ALSO

11965 *feof()*, *ferror()*, *fopen()*, the Base Definitions volume of IEEE Std 1003.1-200x, `<stdio.h>`,  
11966 `<wchar.h>`

#### 11967 CHANGE HISTORY

11968 First released in Issue 4. Derived from the MSE working draft.

#### 11969 Issue 5

11970 The Optional Header (OH) marking is removed from `<stdio.h>`.

11971 Large File Summit extensions are added.

#### 11972 Issue 6

11973 Extensions beyond the ISO C standard are now marked.

11974 The following new requirements on POSIX implementations derive from alignment with the  
11975 Single UNIX Specification:

- 11976 • The [EIO] and [EOVERFLOW] mandatory error conditions are added.
- 11977 • The [ENOMEM] and [ENXIO] optional error conditions are added.

11978 **NAME**11979 `fgetws` — get a wide-character string from a stream11980 **SYNOPSIS**11981 `#include <stdio.h>`11982 `#include <wchar.h>`11983 `wchar_t *fgetws(wchar_t *restrict ws, int n,`11984 `FILE *restrict stream);`11985 **DESCRIPTION**

11986 **CX** The functionality described on this reference page is aligned with the ISO C standard. Any  
 11987 conflict between the requirements described here and the ISO C standard is unintentional. This  
 11988 volume of IEEE Std 1003.1-200x defers to the ISO C standard.

11989 The `fgetws()` function shall read characters from the *stream*, convert these to the corresponding  
 11990 wide-character codes, place them in the `wchar_t` array pointed to by *ws*, until *n*−1 characters are  
 11991 read, or a <newline> is read, converted, and transferred to *ws*, or an end-of-file condition is  
 11992 encountered. The wide-character string, *ws*, shall then be terminated with a null wide-character  
 11993 code.

11994 If an error occurs, the resulting value of the file position indicator for the stream is unspecified.

11995 **CX** The `fgetws()` function may mark the `st_atime` field of the file associated with *stream* for update.  
 11996 The `st_atime` field shall be marked for update by the first successful execution of `fgetc()`, `fgets()`,  
 11997 `fgetwc()`, `fgetws()`, `fread()`, `fscanf()`, `getc()`, `getchar()`, `gets()`, or `scanf()` using *stream* that returns  
 11998 data not supplied by a prior call to `ungetc()` or `ungetwc()`.

11999 **RETURN VALUE**

12000 Upon successful completion, `fgetws()` shall return *ws*. If the stream is at end-of-file, the end-of-  
 12001 file indicator for the stream shall be set and `fgetws()` shall return a null pointer. If a read error  
 12002 **CX** occurs, the error indicator for the stream shall be set, `fgetws()` shall return a null pointer, and  
 12003 shall set `errno` to indicate the error.

12004 **ERRORS**12005 Refer to `fgetwc()`.12006 **EXAMPLES**

12007 None.

12008 **APPLICATION USAGE**

12009 None.

12010 **RATIONALE**

12011 None.

12012 **FUTURE DIRECTIONS**

12013 None.

12014 **SEE ALSO**12015 `fopen()`, `fread()`, the Base Definitions volume of IEEE Std 1003.1-200x, `<stdio.h>`, `<wchar.h>`12016 **CHANGE HISTORY**

12017 First released in Issue 4. Derived from the MSE working draft.

12018 **Issue 5**12019 The Optional Header (OH) marking is removed from `<stdio.h>`.

12020 **Issue 6**

12021 Extensions beyond the ISO C standard are now marked.

12022 The prototype for *fgetws()* is changed for alignment with the ISO/IEC 9899:1999 standard.

12023 **NAME**12024 `fileno` — map a stream pointer to a file descriptor12025 **SYNOPSIS**12026 `cx` `#include <stdio.h>`12027 `int fileno(FILE *stream);`

12028

12029 **DESCRIPTION**12030 The `fileno()` function shall return the integer file descriptor associated with the stream pointed to  
12031 by `stream`.12032 **RETURN VALUE**12033 Upon successful completion, `fileno()` shall return the integer value of the file descriptor  
12034 associated with `stream`. Otherwise, the value `-1` shall be returned and `errno` set to indicate the  
12035 error.12036 **ERRORS**12037 The `fileno()` function may fail if:12038 [EBADF] The `stream` argument is not a valid stream.12039 **EXAMPLES**

12040 None.

12041 **APPLICATION USAGE**

12042 None.

12043 **RATIONALE**12044 Without some specification of which file descriptors are associated with these streams, it is  
12045 impossible for an application to set up the streams for another application it starts with `fork()`  
12046 and `exec`. In particular, it would not be possible to write a portable version of the `sh` command  
12047 interpreter (although there may be other constraints that would prevent that portability).12048 **FUTURE DIRECTIONS**

12049 None.

12050 **SEE ALSO**12051 `fdopen()`, `fopen()`, `stdin`, the Base Definitions volume of IEEE Std 1003.1-200x, `<stdio.h>`, Section  
12052 2.5.1 (on page 485)12053 **CHANGE HISTORY**

12054 First released in Issue 1. Derived from Issue 1 of the SVID.

12055 **Issue 6**12056 The following new requirements on POSIX implementations derive from alignment with the  
12057 Single UNIX Specification:

- 12058
- The [EBADF] optional error condition is added.

12059 **NAME**

12060 flockfile, ftrylockfile, funlockfile — stdio locking functions

12061 **SYNOPSIS**

12062 TSF #include &lt;stdio.h&gt;

```
12063 void flockfile(FILE *file);
12064 int ftrylockfile(FILE *file);
12065 void funlockfile(FILE *file);
12066
```

12067 **DESCRIPTION**

12068 These functions shall provide for explicit application-level locking of stdio (**FILE** \*) objects. |  
12069 These functions can be used by a thread to delineate a sequence of I/O statements that are |  
12070 executed as a unit.

12071 The *flockfile()* function shall acquire for a thread ownership of a (**FILE** \*) object. |

12072 The *ftrylockfile()* function shall acquire for a thread ownership of a (**FILE** \*) object if the object is |  
12073 available; *ftrylockfile()* is a non-blocking version of *flockfile()*.

12074 The *funlockfile()* function shall relinquish the ownership granted to the thread. The behavior is |  
12075 undefined if a thread other than the current owner calls the *funlockfile()* function.

12076 The functions shall behave as if there is a lock count associated with each (**FILE** \*) object. This |  
12077 count is implicitly initialized to zero when the (**FILE** \*) object is created. The (**FILE** \*) object is |  
12078 unlocked when the count is zero. When the count is positive, a single thread owns the (**FILE** \*) |  
12079 object. When the *flockfile()* function is called, if the count is zero or if the count is positive and |  
12080 the caller owns the (**FILE** \*) object, the count shall be incremented. Otherwise, the calling thread |  
12081 shall be suspended, waiting for the count to return to zero. Each call to *funlockfile()* shall |  
12082 decrement the count. This allows matching calls to *flockfile()* (or successful calls to *ftrylockfile()*) |  
12083 and *funlockfile()* to be nested.

12084 All functions that reference (**FILE** \*) objects shall behave as if they use *flockfile()* and *funlockfile()* |  
12085 internally to obtain ownership of these (**FILE** \*) objects.

12086 **RETURN VALUE**

12087 None for *flockfile()* and *funlockfile()*. The *ftrylockfile()* function shall return zero for success and |  
12088 non-zero to indicate that the lock cannot be acquired.

12089 **ERRORS**

12090 No errors are defined.

12091 **EXAMPLES**

12092 None.

12093 **APPLICATION USAGE**

12094 Applications using these functions may be subject to priority inversion, as discussed in the Base |  
12095 Definitions volume of IEEE Std 1003.1-200x, Section 3.285, Priority Inversion.

12096 **RATIONALE**

12097 The *flockfile()* and *funlockfile()* functions provide an orthogonal mutual exclusion lock for each |  
12098 **FILE**. The *ftrylockfile()* function provides a non-blocking attempt to acquire a file lock, |  
12099 analogous to *pthread\_mutex\_trylock()*.

12100 These locks behave as if they are the same as those used internally by *stdio* for thread-safety. |  
12101 This both provides thread-safety of these functions without requiring a second level of internal |  
12102 locking and allows functions in *stdio* to be implemented in terms of other *stdio* functions.

12103 Application writers and implementors should be aware that there are potential deadlock  
12104 problems on **FILE** objects. For example, the line-buffered flushing semantics of *stdio* (requested  
12105 via `{_IOLBF}`) require that certain input operations sometimes cause the buffered contents of  
12106 implementation-defined line-buffered output streams to be flushed. If two threads each hold the  
12107 lock on the other's **FILE**, deadlock ensues. This type of deadlock can be avoided by acquiring  
12108 **FILE** locks in a consistent order. In particular, the line-buffered output stream deadlock can  
12109 typically be avoided by acquiring locks on input streams before locks on output streams if a  
12110 thread would be acquiring both.

12111 In summary, threads sharing *stdio* streams with other threads can use *flockfile()* and *funlockfile()*  
12112 to cause sequences of I/O performed by a single thread to be kept bundled. The only case where  
12113 the use of *flockfile()* and *funlockfile()* is required is to provide a scope protecting uses of the  
12114 `*_unlocked()` functions/macros. This moves the cost/performance tradeoff to the optimal point.

#### 12115 **FUTURE DIRECTIONS**

12116 None.

#### 12117 **SEE ALSO**

12118 *getc\_unlocked()*, *putc\_unlocked()*, the Base Definitions volume of IEEE Std 1003.1-200x, `<stdio.h>`

#### 12119 **CHANGE HISTORY**

12120 First released in Issue 5. Included for alignment with the POSIX Threads Extension.

#### 12121 **Issue 6**

12122 These functions are marked as part of the Thread-Safe Functions option.

## 12123 NAME

12124 floor, floorf, floorl — floor function

## 12125 SYNOPSIS

12126 #include &lt;math.h&gt;

12127 double floor(double x);

12128 float floorf(float x);

12129 long double floorl(long double x);

## 12130 DESCRIPTION

12131 **CX** The functionality described on this reference page is aligned with the ISO C standard. Any  
 12132 conflict between the requirements described here and the ISO C standard is unintentional. This  
 12133 volume of IEEE Std 1003.1-200x defers to the ISO C standard.

12134 These functions shall compute the largest integral value not greater than *x*.

12135 An application wishing to check for error situations should set *errno* to zero and call  
 12136 *feclearexcept*(FE\_ALL\_EXCEPT) before calling these functions. On return, if *errno* is non-zero or  
 12137 *fetestexcept*(FE\_INVALID | FE\_DIVBYZERO | FE\_OVERFLOW | FE\_UNDERFLOW) is non-  
 12138 zero, an error has occurred.

## 12139 RETURN VALUE

12140 Upon successful completion, these functions shall return the largest integral value not greater  
 12141 than *x*, expressed as a **double**, **float**, or **long double**, as appropriate for the return type of the  
 12142 function.

12143 **MX** If *x* is NaN, a NaN shall be returned.12144 If *x* is  $\pm 0$  or  $\pm \text{Inf}$ , *x* shall be returned.

12145 **XSI** If the correct value would cause overflow, a range error shall occur and *floor*(*x*), *floorf*(*x*), and  
 12146 *floorl*(*x*) shall return the value of the macro HUGE\_VAL, HUGE\_VALF, and HUGE\_VALL,  
 12147 respectively.

## 12148 ERRORS

12149 These functions shall fail if:

12150 **XSI** Range Error The result would cause an overflow.

12151 If the integer expression (math\_errhandling & MATH\_ERRNO) is non-zero, |  
 12152 then *errno* shall be set to [ERANGE]. If the integer expression |  
 12153 (math\_errhandling & MATH\_ERREXCEPT) is non-zero, then the overflow |  
 12154 floating-point exception shall be raised. |

## 12155 EXAMPLES

12156 None.

## 12157 APPLICATION USAGE

12158 The integral value returned by these functions might not be expressible as an **int** or **long**. The  
 12159 return value should be tested before assigning it to an integer type to avoid the undefined results  
 12160 of an integer overflow.

12161 The *floor*(*x*) function can only overflow when the floating-point representation has  
 12162 DBL\_MANT\_DIG > DBL\_MAX\_EXP.

12163 On error, the expressions (math\_errhandling & MATH\_ERRNO) and (math\_errhandling &  
 12164 MATH\_ERREXCEPT) are independent of each other, but at least one of them must be non-zero.



12165 **RATIONALE**

12166 None.

12167 **FUTURE DIRECTIONS**

12168 None.

12169 **SEE ALSO**12170 *ceil()*, *feclearexcept()*, *fetestexcept()*, *isnan()*, the Base Definitions volume of IEEE Std 1003.1-200x, |

12171 Section 4.18, Treatment of Error Conditions for Mathematical Functions, &lt;math.h&gt; |

12172 **CHANGE HISTORY**

12173 First released in Issue 1. Derived from Issue 1 of the SVID.

12174 **Issue 5**12175 The DESCRIPTION is updated to indicate how an application should check for an error. This  
12176 text was previously published in the APPLICATION USAGE section.12177 **Issue 6**12178 The *floorf()* and *floorl()* functions are added for alignment with the ISO/IEC 9899:1999 standard.12179 The DESCRIPTION, RETURN VALUE, ERRORS, and APPLICATION USAGE sections are  
12180 revised to align with the ISO/IEC 9899:1999 standard.12181 IEC 60559:1989 standard floating-point extensions over the ISO/IEC 9899:1999 standard are  
12182 marked.

## 12183 NAME

12184 fma, fmaf, fmal — floating-point multiply-add

## 12185 SYNOPSIS

12186 #include &lt;math.h&gt;

12187 double fma(double x, double y, double z);

12188 float fmaf(float x, float y, float z);

12189 long double fmal(long double x, long double y, long double z);

## 12190 DESCRIPTION

12191 **CX** The functionality described on this reference page is aligned with the ISO C standard. Any  
 12192 conflict between the requirements described here and the ISO C standard is unintentional. This  
 12193 volume of IEEE Std 1003.1-200x defers to the ISO C standard.

12194 These functions shall compute  $(x * y) + z$ , rounded as one ternary operation: they shall compute  
 12195 the value (as if) to infinite precision and round once to the result format, according to the  
 12196 rounding mode characterized by the value of FLT\_ROUNDS.

12197 An application wishing to check for error situations should set *errno* to zero and call  
 12198 *feclearexcept*(FE\_ALL\_EXCEPT) before calling these functions. On return, if *errno* is non-zero or  
 12199 *fetestexcept*(FE\_INVALID | FE\_DIVBYZERO | FE\_OVERFLOW | FE\_UNDERFLOW) is non-  
 12200 zero, an error has occurred.

## 12201 RETURN VALUE

12202 Upon successful completion, these functions shall return  $(x * y) + z$ , rounded as one ternary  
 12203 operation.

12204 **MX** If *x* or *y* are NaN, a NaN shall be returned.

12205 If *x* multiplied by *y* is an exact infinity and *z* is also an infinity but with the opposite sign, a  
 12206 domain error shall occur, and either a NaN (if supported), or an implementation-defined value  
 12207 shall be returned.

12208 If one of *x* and *y* is infinite, the other is zero, and *z* is not a NaN, a domain error shall occur, and  
 12209 either a NaN (if supported), or an implementation-defined value shall be returned.

12210 If one of *x* and *y* is infinite, the other is zero, and *z* is a NaN, a NaN shall be returned and a  
 12211 domain error may occur.

12212 If  $x*y$  is not  $0*Inf$  nor  $Inf*0$  and *z* is a NaN, a NaN shall be returned.

## 12213 ERRORS

12214 These functions shall fail if:

12215 **MX** Domain Error The value of  $x*y+z$  is invalid, or the value  $x*y$  is invalid and *z* is not a NaN.

12216 If the integer expression (math\_errhandling & MATH\_ERRNO) is non-zero, |  
 12217 then *errno* shall be set to [EDOM]. If the integer expression (math\_errhandling |  
 12218 & MATH\_ERREXCEPT) is non-zero, then the invalid floating-point exception |  
 12219 shall be raised. |

12220 **MX** Range Error The result overflows.

12221 If the integer expression (math\_errhandling & MATH\_ERRNO) is non-zero, |  
 12222 then *errno* shall be set to [ERANGE]. If the integer expression |  
 12223 (math\_errhandling & MATH\_ERREXCEPT) is non-zero, then the overflow |  
 12224 floating-point exception shall be raised. |

12225 These functions may fail if:

12226	MX	<b>Domain Error</b>	The value $x*y$ is invalid and $z$ is a NaN.
12227			If the integer expression ( <code>math_errhandling &amp; MATH_ERRNO</code> ) is non-zero,
12228			then <i>errno</i> shall be set to [EDOM]. If the integer expression ( <code>math_errhandling</code>
12229			& <code>MATH_ERREXCEPT</code> ) is non-zero, then the invalid floating-point exception
12230			shall be raised.
12231	MX	<b>Range Error</b>	The result underflows.
12232			If the integer expression ( <code>math_errhandling &amp; MATH_ERRNO</code> ) is non-zero,
12233			then <i>errno</i> shall be set to [ERANGE]. If the integer expression
12234			( <code>math_errhandling &amp; MATH_ERREXCEPT</code> ) is non-zero, then the underflow
12235			floating-point exception shall be raised.
12236		<b>EXAMPLES</b>	
12237			None.
12238		<b>APPLICATION USAGE</b>	
12239			On error, the expressions ( <code>math_errhandling &amp; MATH_ERRNO</code> ) and ( <code>math_errhandling &amp;</code>
12240			<code>MATH_ERREXCEPT</code> ) are independent of each other, but at least one of them must be non-zero.
12241		<b>RATIONALE</b>	
12242			In many cases, clever use of floating ( <i>fused</i> ) multiply-add leads to much improved code; but its
12243			unexpected use by the compiler can undermine carefully written code. The <code>FP_CONTRACT</code>
12244			macro can be used to disallow use of floating multiply-add; and the <i>fma()</i> function guarantees
12245			its use where desired. Many current machines provide hardware floating multiply-add
12246			instructions; software implementation can be used for others.
12247		<b>FUTURE DIRECTIONS</b>	
12248			None.
12249		<b>SEE ALSO</b>	
12250			<i>feclearexcept()</i> , <i>fetestexcept()</i> , the Base Definitions volume of IEEE Std 1003.1-200x, Section 4.18,
12251			Treatment of Error Conditions for Mathematical Functions, < <b>math.h</b> >
12252		<b>CHANGE HISTORY</b>	
12253			First released in Issue 6. Derived from the ISO/IEC 9899:1999 standard.

12254 **NAME**

12255 fmax, fmaxf, fmaxl — determine maximum numeric value of two floating-point numbers

12256 **SYNOPSIS**

12257 #include <math.h>

12258 double fmax(double x, double y);

12259 float fmaxf(float x, float y);

12260 long double fmaxl(long double x, long double y);

12261 **DESCRIPTION**

12262 **CX** The functionality described on this reference page is aligned with the ISO C standard. Any  
12263 conflict between the requirements described here and the ISO C standard is unintentional. This  
12264 volume of IEEE Std 1003.1-200x defers to the ISO C standard.

12265 These functions shall determine the maximum numeric value of their arguments. NaN  
12266 arguments shall be treated as missing data: if one argument is a NaN and the other numeric,  
12267 then these functions shall choose the numeric value.

12268 **RETURN VALUE**

12269 Upon successful completion, these functions shall return the maximum numeric value of their  
12270 arguments.

12271 If just one argument is a NaN, the other argument shall be returned.

12272 **MX** If *x* and *y* are NaN, a NaN shall be returned.

12273 **ERRORS**

12274 No errors are defined.

12275 **EXAMPLES**

12276 None.

12277 **APPLICATION USAGE**

12278 None.

12279 **RATIONALE**

12280 None.

12281 **FUTURE DIRECTIONS**

12282 None.

12283 **SEE ALSO**

12284 *fdim()*, *fmin()*, the Base Definitions volume of IEEE Std 1003.1-200x, <math.h>

12285 **CHANGE HISTORY**

12286 First released in Issue 6. Derived from the ISO/IEC 9899:1999 standard.

12287 **NAME**

12288 fmin, fminf, fminl — determine minimum numeric value of two floating-point numbers

12289 **SYNOPSIS**

12290 #include <math.h>

12291 double fmin(double x, double y);

12292 float fminf(float x, float y);

12293 long double fminl(long double x, long double y);

12294 **DESCRIPTION**

12295 **CX** The functionality described on this reference page is aligned with the ISO C standard. Any  
12296 conflict between the requirements described here and the ISO C standard is unintentional. This  
12297 volume of IEEE Std 1003.1-200x defers to the ISO C standard.

12298 These functions shall determine the minimum numeric value of their arguments. NaN  
12299 arguments shall be treated as missing data: if one argument is a NaN and the other numeric,  
12300 then these functions shall choose the numeric value.

12301 **RETURN VALUE**

12302 Upon successful completion, these functions shall return the minimum numeric value of their  
12303 arguments.

12304 If just one argument is a NaN, the other argument shall be returned.

12305 **MX** If *x* and *y* are NaN, a NaN shall be returned.

12306 **ERRORS**

12307 No errors are defined.

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 zero and call  
 12333 *feclearexcept*(FE\_ALL\_EXCEPT) before calling these functions. On return, if *errno* is non-zero or  
 12334 *fetestexcept*(FE\_INVALID | FE\_DIVBYZERO | FE\_OVERFLOW | FE\_UNDERFLOW) is non-  
 12335 zero, an error has occurred.

12336 **RETURN VALUE**

12337 These functions shall return the value  $x-i*y$ , for some integer  $i$  such that, if  $y$  is non-zero, the  
 12338 result has the same sign as  $x$  and magnitude less than the magnitude of  $y$ .

12339 If the correct value would cause underflow, and is not representable, a range error may occur,  
 12340 **MX** and either 0.0 (if supported), or an implementation-defined value shall be returned.

12341 **MX** If  $x$  or  $y$  is NaN, a NaN shall be returned.

12342 If  $y$  is zero, a domain error shall occur, and either a NaN (if supported), or an implementation-  
 12343 defined value shall be returned.

12344 If  $x$  is infinite, a domain error shall occur, and either a NaN (if supported), or an  
 12345 implementation-defined value shall be returned.

12346 If  $x$  is  $\pm 0$  and  $y$  is not zero,  $\pm 0$  shall be returned.12347 If  $x$  is not infinite and  $y$  is  $\pm \text{Inf}$ ,  $x$  shall be returned.

12348 If the correct value would cause underflow, and is representable, a range error may occur and  
 12349 the correct value shall be returned.

12350 **ERRORS**

12351 These functions shall fail if:

12352 **MX** **Domain Error** The  $x$  argument is infinite or  $y$  is zero.

12353 If the integer expression (math\_errhandling & MATH\_ERRNO) is non-zero, |  
 12354 then *errno* shall be set to [EDOM]. If the integer expression (math\_errhandling |  
 12355 & MATH\_ERREXCEPT) is non-zero, then the invalid floating-point exception |  
 12356 shall be raised. |

12357 These functions may fail if:

12358 **Range Error** The result underflows.

12359 If the integer expression (math\_errhandling & MATH\_ERRNO) is non-zero, |  
 12360 then *errno* shall be set to [ERANGE]. If the integer expression |  
 12361 (math\_errhandling & MATH\_ERREXCEPT) is non-zero, then the underflow |  
 12362 floating-point exception shall be raised. |

12363 **EXAMPLES**

12364           None.

12365 **APPLICATION USAGE**

12366           On error, the expressions (math\_errhandling & MATH\_ERRNO) and (math\_errhandling &  
12367           MATH\_ERREXCEPT) are independent of each other, but at least one of them must be non-zero.

12368 **RATIONALE**

12369           None.

12370 **FUTURE DIRECTIONS**

12371           None.

12372 **SEE ALSO**

12373           *feclearexcept()*, *fetestexcept()*, *isnan()*, the Base Definitions volume of IEEE Std 1003.1-200x, |  
12374           Section 4.18, Treatment of Error Conditions for Mathematical Functions, <math.h> |

12375 **CHANGE HISTORY**

12376           First released in Issue 1. Derived from Issue 1 of the SVID.

12377 **Issue 5**

12378           The DESCRIPTION is updated to indicate how an application should check for an error. This  
12379           text was previously published in the APPLICATION USAGE section.

12380 **Issue 6**12381           The behavior for when the *y* argument is zero is now defined.

12382           The *fmodf()* and *fmodl()* functions are added for alignment with the ISO/IEC 9899:1999  
12383           standard.

12384           The DESCRIPTION, RETURN VALUE, ERRORS, and APPLICATION USAGE sections are  
12385           revised to align with the ISO/IEC 9899:1999 standard.

12386           IEC 60559:1989 standard floating-point extensions over the ISO/IEC 9899:1999 standard are  
12387           marked.

12388 **NAME**12389 `fmtmsg` — display a message in the specified format on standard error and/or a system console12390 **SYNOPSIS**

```
12391 xSI #include <fmtmsg.h>
12392
12392 int fmtmsg(long classification, const char *label, int severity,
12393           const char *text, const char *action, const char *tag);
12394
```

12395 **DESCRIPTION**12396 The `fmtmsg()` function shall display messages in a specified format instead of the traditional `printf()` function. 1239712398 Based on a message's classification component, `fmtmsg()` shall write a formatted message either to standard error, to the console, or to both. 1239912400 A formatted message consists of up to five components as defined below. The component `classification` is not part of a message displayed to the user, but defines the source of the message and directs the display of the formatted message. 12401 12402

12403 *classification* Contains the sum of identifying values constructed from the constants defined below. Any one identifier from a subclass may be used in combination with a single identifier from a different subclass. Two or more identifiers from the same subclass should not be used together, with the exception of identifiers from the display subclass. (Both display subclass identifiers may be used so that messages can be displayed to both standard error and the system console). 12404 12405 12406 12407 12408 12409

12410 **Major Classifications**

12411 Identifies the source of the condition. Identifiers are: MM\_HARD (hardware), MM\_SOFT (software), and MM\_FIRM (firmware). 12412

12413 **Message Source Subclassifications**

12414 Identifies the type of software in which the problem is detected. Identifiers are: MM\_APPL (application), MM\_UTIL (utility), and MM\_OPYSYS (operating system). 12415 12416

12417 **Display Subclassifications**

12418 Indicates where the message is to be displayed. Identifiers are: MM\_PRINT to display the message on the standard error stream, MM\_CONSOLE to display the message on the system console. One or both identifiers may be used. 12419 12420 12421

12422 **Status Subclassifications**

12423 Indicates whether the application can recover from the condition. Identifiers are: MM\_RECOVER (recoverable) and MM\_NRECOV (non-recoverable). 12424 12425

12426 An additional identifier, MM\_NULLMC, indicates that no classification component is supplied for the message. 12427

12428 *label* Identifies the source of the message. The format is two fields separated by a colon. The first field is up to 10 bytes, the second is up to 14 bytes. 1242912430 *severity* Indicates the seriousness of the condition. Identifiers for the levels of *severity* are: 12431



12432		MM_HALT	Indicates that the application has encountered a severe fault and is halting. Produces the string "HALT".
12433			
12434		MM_ERROR	Indicates that the application has detected a fault. Produces the string "ERROR".
12435			
12436		MM_WARNING	Indicates a condition that is out of the ordinary, that might be a problem, and should be watched. Produces the string "WARNING".
12437			
12438			
12439		MM_INFO	Provides information about a condition that is not in error. Produces the string "INFO".
12440			
12441		MM_NOSEV	Indicates that no severity level is supplied for the message.
12442	<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.
12443			
12444			
12445	<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.
12446			
12447			
12448	<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".
12449			
12450			
12451			The <i>MSGVERB</i> environment variable (for message verbosity) shall determine for <i>fmtmsg()</i>
12452			which message components it is to select when writing messages to standard error. The value of
12453			<i>MSGVERB</i> shall be a colon-separated list of optional keywords. Valid <i>keywords</i> are: <i>label</i> , <i>severity</i> ,
12454			<i>text</i> , <i>action</i> , and <i>tag</i> . If <i>MSGVERB</i> contains a keyword for a component and the component's
12455			value is not the component's null value, <i>fmtmsg()</i> shall include that component in the message
12456			when writing the message to standard error. If <i>MSGVERB</i> does not include a keyword for a
12457			message component, that component is shall not be included in the display of the message. The
12458			keywords may appear in any order. If <i>MSGVERB</i> is not defined, if its value is the null string, if
12459			its value is not of the correct format, or if it contains keywords other than the valid ones listed
12460			above, <i>fmtmsg()</i> shall select all components.
12461			<i>MSGVERB</i> shall determine which components are selected for display to standard error. All
12462			message components shall be included in console messages.
12463	<b>RETURN VALUE</b>		
12464			The <i>fmtmsg()</i> function shall return one of the following values:
12465	MM_OK		The function succeeded.
12466	MM_NOTOK		The function failed completely.
12467	MM_NOMSG		The function was unable to generate a message on standard error, but
12468			otherwise succeeded.
12469	MM_NOCON		The function was unable to generate a console message, but otherwise
12470			succeeded.
12471	<b>ERRORS</b>		
12472			None.

12473 **EXAMPLES**

12474 1. The following example of *fmtmsg()*:

```
12475     fmtmsg(MM_PRINT, "XSI:cat", MM_ERROR, "illegal option",  
12476     "refer to cat in user's reference manual", "XSI:cat:001")
```

12477 produces a complete message in the specified message format:

```
12478     XSI:cat: ERROR: illegal option  
12479     TO FIX: refer to cat in user's reference manual XSI:cat:001
```

12480 2. When the environment variable *MSGVERB* is set as follows:

```
12481     MSGVERB=severity:text:action
```

12482 and Example 1 is used, *fmtmsg()* produces:

```
12483     ERROR: illegal option  
12484     TO FIX: refer to cat in user's reference manual
```

12485 **APPLICATION USAGE**

12486 One or more message components may be systematically omitted from messages generated by  
12487 an application by using the null value of the argument for that component.

12488 **RATIONALE**

12489 None.

12490 **FUTURE DIRECTIONS**

12491 None.

12492 **SEE ALSO**

12493 *printf()*, the Base Definitions volume of IEEE Std 1003.1-200x, <*fmtmsg.h*>

12494 **CHANGE HISTORY**

12495 First released in Issue 4, Version 2.

12496 **Issue 5**

12497 Moved from X/OPEN UNIX extension to BASE.

12498 **NAME**

12499 fnmatch — match a filename or a pathname |

12500 **SYNOPSIS**

12501 #include &lt;fnmatch.h&gt;

12502 int fnmatch(const char \**pattern*, const char \**string*, int *flags*);12503 **DESCRIPTION**

12504 The *fnmatch()* function shall match patterns as described in the Shell and Utilities volume of  
 12505 IEEE Std 1003.1-200x, Section 2.13.1, Patterns Matching a Single Character, and Section 2.13.2,  
 12506 Patterns Matching Multiple Characters. It checks the string specified by the *string* argument to  
 12507 see if it matches the pattern specified by the *pattern* argument.

12508 The *flags* argument shall modify the interpretation of *pattern* and *string*. It is the bitwise-inclusive  
 12509 OR of zero or more of the flags defined in <fnmatch.h>. If the FNM\_PATHNAME flag is set in  
 12510 *flags*, then a slash character ( '/' ) in *string* shall be explicitly matched by a slash in *pattern*; it shall  
 12511 not be matched by either the asterisk or question-mark special characters, nor by a bracket  
 12512 expression. If the FNM\_PATHNAME flag is not set, the slash character shall be treated as an  
 12513 ordinary character. |

12514 If FNM\_NOESCAPE is not set in *flags*, a backslash character ( '\ ' ) in *pattern* followed by any  
 12515 other character shall match that second character in *string*. In particular, "\\\" shall match a  
 12516 backslash in *string*. If FNM\_NOESCAPE is set, a backslash character shall be treated as an  
 12517 ordinary character.

12518 If FNM\_PERIOD is set in *flags*, then a leading period ( '.' ) in *string* shall match a period in  
 12519 *pattern*; as described by rule 2 in the Shell and Utilities volume of IEEE Std 1003.1-200x, Section  
 12520 2.13.3, Patterns Used for Filename Expansion where the location of “leading” is indicated by the  
 12521 value of FNM\_PATHNAME:

- 12522 • If FNM\_PATHNAME is set, a period is “leading” if it is the first character in *string* or if it  
 12523 immediately follows a slash.

- 12524 • If FNM\_PATHNAME is not set, a period is “leading” only if it is the first character of *string*.

12525 If FNM\_PERIOD is not set, then no special restrictions are placed on matching a period.

12526 **RETURN VALUE**

12527 If *string* matches the pattern specified by *pattern*, then *fnmatch()* shall return 0. If there is no  
 12528 match, *fnmatch()* shall return FNM\_NOMATCH, which is defined in <fnmatch.h>. If an error  
 12529 occurs, *fnmatch()* shall return another non-zero value. |

12530 **ERRORS**

12531 No errors are defined.

12532 **EXAMPLES**

12533 None.

12534 **APPLICATION USAGE**

12535 The *fnmatch()* function has two major uses. It could be used by an application or utility that  
 12536 needs to read a directory and apply a pattern against each entry. The *find* utility is an example of  
 12537 this. It can also be used by the *pax* utility to process its *pattern* operands, or by applications that  
 12538 need to match strings in a similar manner.

12539 The name *fnmatch()* is intended to imply *filename* match, rather than *pathname* match. The default  
 12540 action of this function is to match filenames, rather than pathnames, since it gives no special  
 12541 significance to the slash character. With the FNM\_PATHNAME flag, *fnmatch()* does match  
 12542 pathnames, but without tilde expansion, parameter expansion, or special treatment for a period |

12543 at the beginning of a filename.

12544 **RATIONALE**

12545 This function replaced the REG\_FILENAME flag of *regcomp()* in early proposals of this volume  
12546 of IEEE Std 1003.1-200x. It provides virtually the same functionality as the *regcomp()* and  
12547 *regexec()* functions using the REG\_FILENAME and REG\_FSLASH flags (the REG\_FSLASH flag  
12548 was proposed for *regcomp()*, and would have had the opposite effect from FNM\_PATHNAME),  
12549 but with a simpler function and less system overhead.

12550 **FUTURE DIRECTIONS**

12551 None.

12552 **SEE ALSO**

12553 *glob()*, *wordexp()*, the Base Definitions volume of IEEE Std 1003.1-200x, **<fnmatch.h>**, the Shell  
12554 and Utilities volume of IEEE Std 1003.1-200x

12555 **CHANGE HISTORY**

12556 First released in Issue 4. Derived from the ISO POSIX-2 standard.

12557 **Issue 5**

12558 Moved from POSIX2 C-language Binding to BASE.

## 12559 NAME

12560 fopen — open a stream

## 12561 SYNOPSIS

12562 #include &lt;stdio.h&gt;

12563 FILE \*fopen(const char \*restrict filename, const char \*restrict mode);

## 12564 DESCRIPTION

12565 CX The functionality described on this reference page is aligned with the ISO C standard. Any  
 12566 conflict between the requirements described here and the ISO C standard is unintentional. This  
 12567 volume of IEEE Std 1003.1-200x defers to the ISO C standard.

12568 The *fopen()* function shall open the file whose pathname is the string pointed to by *filename*, and  
 12569 associates a stream with it.

12570 The *mode* argument points to a string. If the string is one of the following, the file shall be opened  
 12571 in the indicated mode. Otherwise, the behavior is undefined.

12572 *r* or *rb* Open file for reading.

12573 *w* or *wb* Truncate to zero length or create file for writing.

12574 *a* or *ab* Append; open or create file for writing at end-of-file.

12575 *r+* or *rb+* or *r+b* Open file for update (reading and writing).

12576 *w+* or *wb+* or *w+b* Truncate to zero length or create file for update.

12577 *a+* or *ab+* or *a+b* Append; open or create file for update, writing at end-of-file.

12578 CX The character 'b' shall have no effect, but is allowed for ISO C standard conformance. Opening  
 12579 a file with read mode (*r* as the first character in the *mode* argument) shall fail if the file does not  
 12580 exist or cannot be read.

12581 Opening a file with append mode (*a* as the first character in the *mode* argument) shall cause all  
 12582 subsequent writes to the file to be forced to the then current end-of-file, regardless of intervening  
 12583 calls to *fseek()*.

12584 When a file is opened with update mode ('+' as the second or third character in the *mode*  
 12585 argument), both input and output may be performed on the associated stream. However, the  
 12586 application shall ensure that output is not directly followed by input without an intervening call  
 12587 to *fflush()* or to a file positioning function (*fseek()*, *fsetpos()*, or *rewind()*), and input is not directly  
 12588 followed by output without an intervening call to a file positioning function, unless the input  
 12589 operation encounters end-of-file.

12590 When opened, a stream is fully buffered if and only if it can be determined not to refer to an  
 12591 interactive device. The error and end-of-file indicators for the stream shall be cleared.

12592 CX If *mode* is *w*, *wb*, *a*, *ab*, *w+*, *wb+*, *w+b*, *a+*, *ab+*, or *a+b*, and the file did not previously exist, upon  
 12593 successful completion, *fopen()* function shall mark for update the *st\_atime*, *st\_ctime*, and *st\_mtime*  
 12594 fields of the file and the *st\_ctime* and *st\_mtime* fields of the parent directory.

12595 If *mode* is *w*, *wb*, *w+*, *wb+*, or *w+b*, and the file did previously exist, upon successful completion,  
 12596 *fopen()* shall mark for update the *st\_ctime* and *st\_mtime* fields of the file. The *fopen()* function  
 12597 shall allocate a file descriptor as *open()* does.

12598 XSI After a successful call to the *fopen()* function, the orientation of the stream shall be cleared, the  
 12599 encoding rule shall be cleared, and the associated *mbstate\_t* object shall be set to describe an  
 12600 initial conversion state.

12601 CX The largest value that can be represented correctly in an object of type `off_t` shall be established  
 12602 as the offset maximum in the open file description.

#### 12603 RETURN VALUE

12604 Upon successful completion, `fopen()` shall return a pointer to the object controlling the stream.  
 12605 CX Otherwise, a null pointer shall be returned, and `errno` shall be set to indicate the error.

#### 12606 ERRORS

12607 The `fopen()` function shall fail if:

12608 CX [EACCES] Search permission is denied on a component of the path prefix, or the file  
 12609 exists and the permissions specified by `mode` are denied, or the file does not  
 12610 exist and write permission is denied for the parent directory of the file to be  
 12611 created.

12612 CX [EINTR] A signal was caught during `fopen()`.

12613 CX [EISDIR] The named file is a directory and `mode` requires write access.

12614 CX [ELOOP] A loop exists in symbolic links encountered during resolution of the `path`  
 12615 argument.

12616 CX [EMFILE] {OPEN\_MAX} file descriptors are currently open in the calling process.

12617 CX [ENAMETOOLONG]

12618 The length of the `filename` argument exceeds {PATH\_MAX} or a pathname  
 12619 component is longer than {NAME\_MAX}.

12620 CX [ENFILE] The maximum allowable number of files is currently open in the system.

12621 CX [ENOENT] A component of `filename` does not name an existing file or `filename` is an empty  
 12622 string.

12623 CX [ENOSPC] The directory or file system that would contain the new file cannot be  
 12624 expanded, the file does not exist, and it was to be created.

12625 CX [ENOTDIR] A component of the path prefix is not a directory.

12626 CX [ENXIO] The named file is a character special or block special file, and the device  
 12627 associated with this special file does not exist.

12628 CX [EOVERFLOW] The named file is a regular file and the size of the file cannot be represented  
 12629 correctly in an object of type `off_t`.

12630 CX [EROFS] The named file resides on a read-only file system and `mode` requires write  
 12631 access.

12632 The `fopen()` function may fail if:

12633 CX [EINVAL] The value of the `mode` argument is not valid.

12634 CX [ELOOP] More than {SYMLOOP\_MAX} symbolic links were encountered during  
 12635 resolution of the `path` argument.

12636 CX [EMFILE] {FOPEN\_MAX} streams are currently open in the calling process.

12637 CX [EMFILE] {STREAM\_MAX} streams are currently open in the calling process.

12638 CX [ENAMETOOLONG]

12639 Pathname resolution of a symbolic link produced an intermediate result  
 12640 whose length exceeds {PATH\_MAX}.

- 12641 CX [ENOMEM] Insufficient storage space is available.
- 12642 CX [ETXTBSY] The file is a pure procedure (shared text) file that is being executed and *mode* requires write access.
- 12643

12644 **EXAMPLES**12645 **Opening a File**

12646 The following example tries to open the file named **file** for reading. The *fopen()* function returns  
 12647 a file pointer that is used in subsequent *fgets()* and *fclose()* calls. If the program cannot open the  
 12648 file, it just ignores it.

```
12649 #include <stdio.h>
12650 ...
12651 FILE *fp;
12652 ...
12653 void rgrep(const char *file)
12654 {
12655     ...
12656     if ((fp = fopen(file, "r")) == NULL)
12657         return;
12658     ...
12659 }
```

12660 **APPLICATION USAGE**

12661 None.

12662 **RATIONALE**

12663 None.

12664 **FUTURE DIRECTIONS**

12665 None.

12666 **SEE ALSO**

12667 *fclose()*, *fdopen()*, *freopen()*, the Base Definitions volume of IEEE Std 1003.1-200x, <stdio.h>

12668 **CHANGE HISTORY**

12669 First released in Issue 1. Derived from Issue 1 of the SVID.

12670 **Issue 5**

12671 Large File Summit extensions are added.

12672 **Issue 6**

12673 Extensions beyond the ISO C standard are now marked.

12674 The following new requirements on POSIX implementations derive from alignment with the  
 12675 Single UNIX Specification:

- 12676 • In the DESCRIPTION, text is added to indicate setting of the offset maximum in the open file  
 12677 description. This change is to support large files.
- 12678 • In the ERRORS section, the [Eoverflow] condition is added. This change is to support  
 12679 large files.
- 12680 • The [ELOOP] mandatory error condition is added.
- 12681 • The [EINVAL], [EMFILE], [ENAMETOOLONG], [ENOMEM], and [ETXTBSY] optional error  
 12682 conditions are added.

- 12683 The DESCRIPTION is updated to avoid use of the term “must” for application requirements.
- 12684 The following changes are made for alignment with the ISO/IEC 9899: 1999 standard:
- 12685 • The prototype for *fopen()* is updated.
  - 12686 • The DESCRIPTION is updated to note that if the argument *mode* points to a string other than  
12687 those listed, then the behavior is undefined.
- 12688 The wording of the mandatory [ELOOP] error condition is updated, and a second optional  
12689 [ELOOP] error condition is added.



## 12690 NAME

12691 fork — create a new process

## 12692 SYNOPSIS

12693 #include &lt;unistd.h&gt;

12694 pid\_t fork(void);

## 12695 DESCRIPTION

12696 The *fork()* function shall create a new process. The new process (child process) shall be an exact  
 12697 copy of the calling process (parent process) except as detailed below:

- 12698 • The child process shall have a unique process ID. |
- 12699 • The child process ID also shall not match any active process group ID. |
- 12700 • The child process shall have a different parent process ID, which shall be the process ID of |
- 12701 the calling process. |
- 12702 • The child process shall have its own copy of the parent's file descriptors. Each of the child's |
- 12703 file descriptors shall refer to the same open file description with the corresponding file |
- 12704 descriptor of the parent. |
- 12705 • The child process shall have its own copy of the parent's open directory streams. Each open |
- 12706 directory stream in the child process may share directory stream positioning with the |
- 12707 corresponding directory stream of the parent. |
- 12708 XSI • The child process shall have its own copy of the parent's message catalog descriptors. |
- 12709 • The child process' values of *tms\_utime*, *tms\_stime*, *tms\_cutime*, and *tms\_cstime* shall be set to 0. |
- 12710 • The time left until an alarm clock signal shall be reset to zero, and the alarm, if any, shall be |
- 12711 canceled; see *alarm()*. |
- 12712 XSI • All *semadj* values shall be cleared. |
- 12713 • File locks set by the parent process shall not be inherited by the child process. |
- 12714 • The set of signals pending for the child process shall be initialized to the empty set. |
- 12715 XSI • Interval timers shall be reset in the child process. |
- 12716 SEM • Any semaphores that are open in the parent process shall also be open in the child process. |
- 12717 ML • The child process shall not inherit any address space memory locks established by the parent |
- 12718 process via calls to *mlockall()* or *mlock()*. |
- 12719 MF|SHM • Memory mappings created in the parent shall be retained in the child process. |
- 12720 MAP\_PRIVATE mappings inherited from the parent shall also be MAP\_PRIVATE mappings |
- 12721 in the child, and any modifications to the data in these mappings made by the parent prior to |
- 12722 calling *fork()* shall be visible to the child. Any modifications to the data in MAP\_PRIVATE |
- 12723 mappings made by the parent after *fork()* returns shall be visible only to the parent. |
- 12724 Modifications to the data in MAP\_PRIVATE mappings made by the child shall be visible only |
- 12725 to the child. |
- 12726 PS • For the SCHED\_FIFO and SCHED\_RR scheduling policies, the child process shall inherit the |
- 12727 policy and priority settings of the parent process during a *fork()* function. For other |
- 12728 scheduling policies, the policy and priority settings on *fork()* are implementation-defined. |
- 12729 TMR • Per-process timers created by the parent shall not be inherited by the child process. |
- 12730 MSG • The child process shall have its own copy of the message queue descriptors of the parent. |
- 12731 Each of the message descriptors of the child shall refer to the same open message queue |

12732		description as the corresponding message descriptor of the parent.	
12733	AIO	• No asynchronous input or asynchronous output operations shall be inherited by the child	
12734		process.	
12735		• A process shall be created with a single thread. If a multi-threaded process calls <i>fork()</i> , the	
12736		new process shall contain a replica of the calling thread and its entire address space, possibly	
12737		including the states of mutexes and other resources. Consequently, to avoid errors, the child	
12738		process may only execute async-signal-safe operations until such time as one of the <i>exec</i>	
12739	THR	functions is called. Fork handlers may be established by means of the <i>pthread_atfork()</i>	
12740		function in order to maintain application invariants across <i>fork()</i> calls.	
12741	TRC TRI	• If the Trace option and the Trace Inherit option are both supported:	
12742		If the calling process was being traced in a trace stream that had its inheritance policy set to	
12743		POSIX_TRACE_INHERITED, the child process shall be traced into that trace stream, and the	
12744		child process shall inherit the parent's mapping of trace event names to trace event type	
12745		identifiers. If the trace stream in which the calling process was being traced had its	
12746		inheritance policy set to POSIX_TRACE_CLOSE_FOR_CHILD, the child process shall not be	
12747		traced into that trace stream. The inheritance policy is set by a call to the	
12748		<i>posix_trace_attr_setinherited()</i> function.	
12749	TRC	• If the Trace option is supported, but the Trace Inherit option is not supported:	
12750		The child process shall not be traced into any of the trace streams of its parent process.	
12751	TRC	• If the Trace option is supported, the child process of a trace controller process shall not	
12752		control the trace streams controlled by its parent process.	
12753	CPT	• The initial value of the CPU-time clock of the child process shall be set to zero.	
12754	TCT	• The initial value of the CPU-time clock of the single thread of the child process shall be set to	
12755		zero.	
12756		All other process characteristics defined by IEEE Std 1003.1-200x shall be the same in the parent	
12757		and child processes. The inheritance of process characteristics not defined by	
12758		IEEE Std 1003.1-200x is unspecified by IEEE Std 1003.1-200x.	
12759		After <i>fork()</i> , both the parent and the child processes shall be capable of executing independently	
12760		before either one terminates.	
12761		<b>RETURN VALUE</b>	
12762		Upon successful completion, <i>fork()</i> shall return 0 to the child process and shall return the	
12763		process ID of the child process to the parent process. Both processes shall continue to execute	
12764		from the <i>fork()</i> function. Otherwise, -1 shall be returned to the parent process, no child process	
12765		shall be created, and <i>errno</i> shall be set to indicate the error.	
12766		<b>ERRORS</b>	
12767		The <i>fork()</i> function shall fail if:	
12768	[EAGAIN]	The system lacked the necessary resources to create another process, or the	
12769		system-imposed limit on the total number of processes under execution	
12770		system-wide or by a single user {CHILD_MAX} would be exceeded.	
12771		The <i>fork()</i> function may fail if:	
12772	[ENOMEM]	Insufficient storage space is available.	

12773 **EXAMPLES**

12774 None.

12775 **APPLICATION USAGE**

12776 None.

12777 **RATIONALE**

12778 Many historical implementations have timing windows where a signal sent to a process group  
 12779 (for example, an interactive SIGINT) just prior to or during execution of *fork()* is delivered to the  
 12780 parent following the *fork()* but not to the child because the *fork()* code clears the child's set of  
 12781 pending signals. This volume of IEEE Std 1003.1-200x does not require, or even permit, this  
 12782 behavior. However, it is pragmatic to expect that problems of this nature may continue to exist  
 12783 in implementations that appear to conform to this volume of IEEE Std 1003.1-200x and pass  
 12784 available verification suites. This behavior is only a consequence of the implementation failing to  
 12785 make the interval between signal generation and delivery totally invisible. From the  
 12786 application's perspective, a *fork()* call should appear atomic. A signal that is generated prior to  
 12787 the *fork()* should be delivered prior to the *fork()*. A signal sent to the process group after the  
 12788 *fork()* should be delivered to both parent and child. The implementation may actually initialize  
 12789 internal data structures corresponding to the child's set of pending signals to include signals  
 12790 sent to the process group during the *fork()*. Since the *fork()* call can be considered as atomic  
 12791 from the application's perspective, the set would be initialized as empty and such signals would  
 12792 have arrived after the *fork()*; see also <signal.h>.

12793 One approach that has been suggested to address the problem of signal inheritance across *fork()*  
 12794 is to add an [EINTR] error, which would be returned when a signal is detected during the call.  
 12795 While this is preferable to losing signals, it was not considered an optimal solution. Although it  
 12796 is not recommended for this purpose, such an error would be an allowable extension for an  
 12797 implementation.

12798 The [ENOMEM] error value is reserved for those implementations that detect and distinguish  
 12799 such a condition. This condition occurs when an implementation detects that there is not enough  
 12800 memory to create the process. This is intended to be returned when [EAGAIN] is inappropriate  
 12801 because there can never be enough memory (either primary or secondary storage) to perform the  
 12802 operation. Since *fork()* duplicates an existing process, this must be a condition where there is  
 12803 sufficient memory for one such process, but not for two. Many historical implementations  
 12804 actually return [ENOMEM] due to temporary lack of memory, a case that is not generally  
 12805 distinct from [EAGAIN] from the perspective of a conforming application.

12806 Part of the reason for including the optional error [ENOMEM] is because the SVID specifies it  
 12807 and it should be reserved for the error condition specified there. The condition is not applicable  
 12808 on many implementations.

12809 IEEE Std 1003.1-1988 neglected to require concurrent execution of the parent and child of *fork()*.  
 12810 A system that single-threads processes was clearly not intended and is considered an  
 12811 unacceptable "toy implementation" of this volume of IEEE Std 1003.1-200x. The only objection  
 12812 anticipated to the phrase "executing independently" is testability, but this assertion should be  
 12813 testable. Such tests require that both the parent and child can block on a detectable action of the  
 12814 other, such as a write to a pipe or a signal. An interactive exchange of such actions should be  
 12815 possible for the system to conform to the intent of this volume of IEEE Std 1003.1-200x.

12816 The [EAGAIN] error exists to warn applications that such a condition might occur. Whether it  
 12817 occurs or not is not in any practical sense under the control of the application because the  
 12818 condition is usually a consequence of the user's use of the system, not of the application's code.  
 12819 Thus, no application can or should rely upon its occurrence under any circumstances, nor  
 12820 should the exact semantics of what concept of "user" is used be of concern to the application  
 12821 writer. Validation writers should be cognizant of this limitation.

12822 There are two reasons why POSIX programmers call *fork()*. One reason is to create a new thread  
 12823 of control within the same program (which was originally only possible in POSIX by creating a  
 12824 new process); the other is to create a new process running a different program. In the latter case,  
 12825 the call to *fork()* is soon followed by a call to one of the *exec* functions.

12826 The general problem with making *fork()* work in a multi-threaded world is what to do with all  
 12827 of the threads. There are two alternatives. One is to copy all of the threads into the new process.  
 12828 This causes the programmer or implementation to deal with threads that are suspended on  
 12829 system calls or that might be about to execute system calls that should not be executed in the  
 12830 new process. The other alternative is to copy only the thread that calls *fork()*. This creates the  
 12831 difficulty that the state of process-local resources is usually held in process memory. If a thread  
 12832 that is not calling *fork()* holds a resource, that resource is never released in the child process  
 12833 because the thread whose job it is to release the resource does not exist in the child process.

12834 When a programmer is writing a multi-threaded program, the first described use of *fork()*,  
 12835 creating new threads in the same program, is provided by the *pthread\_create()* function. The  
 12836 *fork()* function is thus used only to run new programs, and the effects of calling functions that  
 12837 require certain resources between the call to *fork()* and the call to an *exec* function are undefined.

12838 The addition of the *forkall()* function to the standard was considered and rejected. The *forkall()*  
 12839 function lets all the threads in the parent be duplicated in the child. This essentially duplicates  
 12840 the state of the parent in the child. This allows threads in the child to continue processing and  
 12841 allows locks and the state to be preserved without explicit *pthread\_atfork()* code. The calling  
 12842 process has to ensure that the threads processing state that is shared between the parent and  
 12843 child (that is, file descriptors or MAP\_SHARED memory) behaves properly after *forkall()*. For  
 12844 example, if a thread is reading a file descriptor in the parent when *forkall()* is called, then two  
 12845 threads (one in the parent and one in the child) are reading the file descriptor after the *forkall()*.  
 12846 If this is not desired behavior, the parent process has to synchronize with such threads before  
 12847 calling *forkall()*.

12848 When *forkall()* is called, threads, other than the calling thread, that are in POSIX System  
 12849 Interfaces functions that can return with an [EINTR] error may have those functions return  
 12850 [EINTR] if the implementation cannot ensure that the function behaves correctly in the parent  
 12851 and child. In particular, *pthread\_cond\_wait()* and *pthread\_cond\_timedwait()* need to return in order  
 12852 to ensure that the condition has not changed. These functions can be awakened by a spurious  
 12853 condition wakeup rather than returning [EINTR].

#### 12854 FUTURE DIRECTIONS

12855 None.

#### 12856 SEE ALSO

12857 *alarm()*, *exec*, *fcntl()*, *posix\_trace\_attr\_getinherited()*, *posix\_trace\_trid\_eventid\_open()*, *semop()*,  
 12858 *signal()*, *times()*, the Base Definitions volume of IEEE Std 1003.1-200x, <sys/types.h>, <unistd.h>

#### 12859 CHANGE HISTORY

12860 First released in Issue 1. Derived from Issue 1 of the SVID.

#### 12861 Issue 5

12862 The DESCRIPTION is changed for alignment with the POSIX Realtime Extension and the POSIX  
 12863 Threads Extension.

#### 12864 Issue 6

12865 The following new requirements on POSIX implementations derive from alignment with the  
 12866 Single UNIX Specification:

- 12867 • The requirement to include <sys/types.h> has been removed. Although <sys/types.h> was  
 12868 required for conforming implementations of previous POSIX specifications, it was not

12869 required for UNIX applications.

12870 The following changes were made to align with the IEEE P1003.1a draft standard:

12871 • The effect of *fork()* on a pending alarm call in the child process is clarified.

12872 The description of CPU-time clock semantics is added for alignment with IEEE Std 1003.1d-1999.

12873 The description of tracing semantics is added for alignment with IEEE Std 1003.1q-2000.

12874 **NAME**

12875 `fpathconf, pathconf` — get configurable pathname variables

12876 **SYNOPSIS**

12877 `#include <unistd.h>`

12878 `long fpathconf(int fildev, int name);`

12879 `long pathconf(const char *path, int name);`

12880 **DESCRIPTION**

12881 The `fpathconf()` and `pathconf()` functions shall determine the current value of a configurable limit  
12882 or option (*variable*) that is associated with a file or directory.

12883 For `pathconf()`, the *path* argument points to the pathname of a file or directory.

12884 For `fpathconf()`, the *fildev* argument is an open file descriptor.

12885 The *name* argument represents the variable to be queried relative to that file or directory.  
12886 Implementations shall support all of the variables listed in the following table and may support  
12887 others. The variables in the following table come from `<limits.h>` or `<unistd.h>` and the  
12888 symbolic constants, defined in `<unistd.h>`, are the corresponding values used for *name*. Support  
12889 for some pathname configuration variables is dependent on implementation options (see  
12890 shading and margin codes in the table below). Where an implementation option is not  
12891 supported, the variable need not be supported.

12892

12893

12894

12895

12896

12897

12898

12899

12900

12901 ADV

12902 ADV

12903 ADV

12904 ADV

12905 ADV

12906

12907

12908

12909

12910

12911

12912

Variable	Value of <i>name</i>	Requirements
{FILESIZEBITS}	_PC_FILESIZEBITS	3, 4
{LINK_MAX}	_PC_LINK_MAX	1
{MAX_CANON}	_PC_MAX_CANON	2
{MAX_INPUT}	_PC_MAX_INPUT	2
{NAME_MAX}	_PC_NAME_MAX	3, 4
{PATH_MAX}	_PC_PATH_MAX	4, 5
{PIPE_BUF}	_PC_PIPE_BUF	6
{POSIX_ALLOC_SIZE_MIN}	_PC_ALLOC_SIZE_MIN	
{POSIX_REC_INCR_XFER_SIZE}	_PC_REC_INCR_XFER_SIZE	
{POSIX_REC_MAX_XFER_SIZE}	_PC_REC_MAX_XFER_SIZE	
{POSIX_REC_MIN_XFER_SIZE}	_PC_REC_MIN_XFER_SIZE	
{POSIX_REC_XFER_ALIGN}	_PC_REC_XFER_ALIGN	
{SYMLINK_MAX}	_PC_SYMLINK_MAX	4, 9
_POSIX_CHOWN_RESTRICTED	_PC_CHOWN_RESTRICTED	7
_POSIX_NO_TRUNC	_PC_NO_TRUNC	3, 4
_POSIX_VDISABLE	_PC_VDISABLE	2
_POSIX_ASYNC_IO	_PC_ASYNC_IO	8
_POSIX_PRIO_IO	_PC_PRIO_IO	8
_POSIX_SYNC_IO	_PC_SYNC_IO	8

12913	<b>Requirements</b>	
12914	1. If <i>path</i> or <i>fildev</i> refers to a directory, the value returned shall apply to the directory itself.	
12915	2. If <i>path</i> or <i>fildev</i> does not refer to a terminal file, it is unspecified whether an implementation	
12916	supports an association of the variable name with the specified file.	
12917	3. If <i>path</i> or <i>fildev</i> refers to a directory, the value returned shall apply to filenames within the	
12918	directory.	
12919	4. If <i>path</i> or <i>fildev</i> does not refer to a directory, it is unspecified whether an implementation	
12920	supports an association of the variable name with the specified file.	
12921	5. If <i>path</i> or <i>fildev</i> refers to a directory, the value returned shall be the maximum length of a	
12922	relative pathname when the specified directory is the working directory.	
12923	6. If <i>path</i> refers to a FIFO, or <i>fildev</i> refers to a pipe or FIFO, the value returned shall apply to	
12924	the referenced object. If <i>path</i> or <i>fildev</i> refers to a directory, the value returned shall apply to	
12925	any FIFO that exists or can be created within the directory. If <i>path</i> or <i>fildev</i> refers to any	
12926	other type of file, it is unspecified whether an implementation supports an association of	
12927	the variable name with the specified file.	
12928	7. If <i>path</i> or <i>fildev</i> refers to a directory, the value returned shall apply to any files, other than	
12929	directories, that exist or can be created within the directory.	
12930	8. If <i>path</i> or <i>fildev</i> refers to a directory, it is unspecified whether an implementation supports	
12931	an association of the variable name with the specified file.	
12932	9. If <i>path</i> or <i>fildev</i> refers to a directory, the value returned shall be the maximum length of the	
12933	string that a symbolic link in that directory can contain.	
12934	<b>RETURN VALUE</b>	
12935	If <i>name</i> is an invalid value, both <i>pathconf()</i> and <i>fpathconf()</i> shall return $-1$ and set <i>errno</i> to	
12936	indicate the error.	
12937	If the variable corresponding to <i>name</i> has no limit for the <i>path</i> or file descriptor, both <i>pathconf()</i>	
12938	and <i>fpathconf()</i> shall return $-1$ without changing <i>errno</i> . If the implementation needs to use <i>path</i>	
12939	to determine the value of <i>name</i> and the implementation does not support the association of <i>name</i>	
12940	with the file specified by <i>path</i> , or if the process did not have appropriate privileges to query the	
12941	file specified by <i>path</i> , or <i>path</i> does not exist, <i>pathconf()</i> shall return $-1$ and set <i>errno</i> to indicate the	
12942	error.	
12943	If the implementation needs to use <i>fildev</i> to determine the value of <i>name</i> and the implementation	
12944	does not support the association of <i>name</i> with the file specified by <i>fildev</i> , or if <i>fildev</i> is an invalid	
12945	file descriptor, <i>fpathconf()</i> shall return $-1$ and set <i>errno</i> to indicate the error.	
12946	Otherwise, <i>pathconf()</i> or <i>fpathconf()</i> shall return the current variable value for the file or	
12947	directory without changing <i>errno</i> . The value returned shall not be more restrictive than the	
12948	corresponding value available to the application when it was compiled with the	
12949	implementation's <code>&lt;limits.h&gt;</code> or <code>&lt;unistd.h&gt;</code> .	
12950	<b>ERRORS</b>	
12951	The <i>pathconf()</i> function shall fail if:	
12952	[EINVAL]           The value of <i>name</i> is not valid.	
12953	[ELOOP]            A loop exists in symbolic links encountered during resolution of the <i>path</i>	
12954	argument.	
12955	The <i>pathconf()</i> function may fail if:	

12956	[EACCES]	Search permission is denied for a component of the path prefix.
12957	[EINVAL]	The implementation does not support an association of the variable <i>name</i> with the specified file.
12958		
12959	[ELOOP]	More than {SYMLOOP_MAX} symbolic links were encountered during resolution of the <i>path</i> argument.
12960		
12961	[ENAMETOOLONG]	
12962		The length of the <i>path</i> argument exceeds {PATH_MAX} or a pathname
12963		component is longer than {NAME_MAX}.
12964	[ENAMETOOLONG]	
12965		As a result of encountering a symbolic link in resolution of the <i>path</i> argument,
12966		the length of the substituted pathname string exceeded {PATH_MAX}.
12967	[ENOENT]	A component of <i>path</i> does not name an existing file or <i>path</i> is an empty string.
12968	[ENOTDIR]	A component of the path prefix is not a directory.
12969		The <i>fpathconf()</i> function shall fail if:
12970	[EINVAL]	The value of <i>name</i> is not valid.
12971		The <i>fpathconf()</i> function may fail if:
12972	[EBADF]	The <i>fdes</i> argument is not a valid file descriptor.
12973	[EINVAL]	The implementation does not support an association of the variable <i>name</i> with the specified file.
12974		
12975	<b>EXAMPLES</b>	
12976		None.
12977	<b>APPLICATION USAGE</b>	
12978		None.
12979	<b>RATIONALE</b>	
12980		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.
12981		
12982		
12983		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 pathname components to 14 bytes, as this would be the value specified in <limits.h> on such a system.
12984		
12985		
12986		
12987		Therefore, various useful values can be queried on any pathname or file descriptor, assuming that the appropriate permissions are in place.
12988		
12989		The value returned for the variable {PATH_MAX} indicates the longest relative pathname 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 pathname crosses into a more restrictive file system.
12990		
12991		
12992		
12993		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.)
12994		
12995		
12996		
12997		Unlike the values returned by <i>sysconf()</i> , the pathname-oriented variables are potentially more volatile and are not guaranteed to remain constant throughout the process' lifetime. For
12998		



12999 example, in between two calls to *pathconf()*, the file system in question may have been  
 13000 unmounted and remounted with different characteristics.

13001 Also note that most of the errors are optional. If one of the variables always has the same value  
 13002 on an implementation, the implementation need not look at *path* or *fildev* to return that value and  
 13003 is, therefore, not required to detect any of the errors except the meaning of [EINVAL] that  
 13004 indicates that the value of *name* is not valid for that variable.

13005 If the value of any of the limits are unspecified (logically infinite), they will not be defined in  
 13006 <limits.h> and the *pathconf()* and *fpathconf()* functions return -1 without changing *errno*. This  
 13007 can be distinguished from the case of giving an unrecognized *name* argument because *errno* is set  
 13008 to [EINVAL] in this case.

13009 Since -1 is a valid return value for the *pathconf()* and *fpathconf()* functions, applications should  
 13010 set *errno* to zero before calling them and check *errno* only if the return value is -1.

13011 For the case of {SYMLINK\_MAX}, since both *pathconf()* and *open()* follow symbolic links, there  
 13012 is no way that *path* or *fildev* could refer to a symbolic link.

### 13013 FUTURE DIRECTIONS

13014 None.

### 13015 SEE ALSO

13016 *confstr()*, *sysconf()*, the Base Definitions volume of IEEE Std 1003.1-200x, <limits.h>, <unistd.h>,  
 13017 the Shell and Utilities volume of IEEE Std 1003.1-200x

### 13018 CHANGE HISTORY

13019 First released in Issue 3.

13020 Entry included for alignment with the POSIX.1-1988 standard.

### 13021 Issue 5

13022 The DESCRIPTION is updated for alignment with the POSIX Realtime Extension.

13023 Large File Summit extensions are added.

### 13024 Issue 6

13025 The following new requirements on POSIX implementations derive from alignment with the  
 13026 Single UNIX Specification:

- 13027 • The DESCRIPTION is updated to include {FILESIZEBITS}.
- 13028 • The [ELOOP] mandatory error condition is added.
- 13029 • A second [ENAMETOOLONG] is added as an optional error condition.

13030 The following changes were made to align with the IEEE P1003.1a draft standard:

- 13031 • The \_PC\_SYMLINK\_MAX entry is added to the table in the DESCRIPTION.

13032 The *pathconf()* variables {POSIX\_ALLOC\_SIZE\_MIN}, {POSIX\_REC\_INCR\_XFER\_SIZE},  
 13033 {POSIX\_REC\_MAX\_XFER\_SIZE}, {POSIX\_REC\_MIN\_XFER\_SIZE},  
 13034 {POSIX\_REC\_XFER\_ALIGN} and their associated names are added for alignment with  
 13035 IEEE Std 1003.1d-1999.

13036 **NAME**

13037 fpclassify — classify real floating type

13038 **SYNOPSIS**

13039 #include &lt;math.h&gt;

13040 int fpclassify(real-floating x);

13041 **DESCRIPTION**

13042 cx The functionality described on this reference page is aligned with the ISO C standard. Any  
13043 conflict between the requirements described here and the ISO C standard is unintentional. This  
13044 volume of IEEE Std 1003.1-200x defers to the ISO C standard.

13045 The *fpclassify()* macro shall classify its argument value as NaN, infinite, normal, subnormal,  
13046 zero, or into another implementation-defined category. First, an argument represented in a  
13047 format wider than its semantic type is converted to its semantic type. Then classification is based  
13048 on the type of the argument.

13049 **RETURN VALUE**

13050 The *fpclassify()* macro shall return the value of the number classification macro appropriate to  
13051 the value of its argument.

13052 **ERRORS**

13053 No errors are defined.

13054 **EXAMPLES**

13055 None.

13056 **APPLICATION USAGE**

13057 None.

13058 **RATIONALE**

13059 None.

13060 **FUTURE DIRECTIONS**

13061 None.

13062 **SEE ALSO**

13063 *isfinite()*, *isinf()*, *isnan()*, *isnormal()*, *signbit()*, the Base Definitions volume of  
13064 IEEE Std 1003.1-200x, <math.h>

13065 **CHANGE HISTORY**

13066 First released in Issue 6. Derived from the ISO/IEC 9899:1999 standard.

13067 **NAME**

13068 fprintf, printf, snprintf, sprintf — print formatted output

13069 **SYNOPSIS**

13070 #include &lt;stdio.h&gt;

13071 int fprintf(FILE \*restrict *stream*, const char \*restrict *format*, ...);13072 int printf(const char \*restrict *format*, ...);13073 int snprintf(char \*restrict *s*, size\_t *n*,13074 const char \*restrict *format*, ...);13075 int sprintf(char \*restrict *s*, const char \*restrict *format*, ...);13076 **DESCRIPTION**

13077 cx The functionality described on this reference page is aligned with the ISO C standard. Any  
 13078 conflict between the requirements described here and the ISO C standard is unintentional. This  
 13079 volume of IEEE Std 1003.1-200x defers to the ISO C standard.

13080 The *fprintf()* function shall place output on the named output *stream*. The *printf()* function shall  
 13081 place output on the standard output stream *stdout*. The *sprintf()* function shall place output  
 13082 followed by the null byte, '\0', in consecutive bytes starting at *s*; it is the user's responsibility  
 13083 to ensure that enough space is available.

13084 The *snprintf()* function shall be equivalent to *sprintf()*, with the addition of the *n* argument  
 13085 which states the size of the buffer referred to by *s*. If *n* is zero, nothing shall be written and *s* may  
 13086 be a null pointer. Otherwise, output bytes beyond the *n*-1st shall be discarded instead of being  
 13087 written to the array, and a null byte is written at the end of the bytes actually written into the  
 13088 array.

13089 If copying takes place between objects that overlap as a result of a call to *sprintf()* or *snprintf()*,  
 13090 the results are undefined.

13091 Each of these functions converts, formats, and prints its arguments under control of the *format*.  
 13092 The *format* is a character string, beginning and ending in its initial shift state, if any. The *format* is  
 13093 composed of zero or more directives: *ordinary characters*, which are simply copied to the output  
 13094 stream, and *conversion specifications*, each of which shall result in the fetching of zero or more  
 13095 arguments. The results are undefined if there are insufficient arguments for the *format*. If the  
 13096 *format* is exhausted while arguments remain, the excess arguments shall be evaluated but are  
 13097 otherwise ignored.

13098 xsi Conversions can be applied to the *n*th argument after the *format* in the argument list, rather than  
 13099 to the next unused argument. In this case, the conversion specifier character % (see below) is  
 13100 replaced by the sequence "%n\$", where *n* is a decimal integer in the range [1,{NL\_ARGMAX}],  
 13101 giving the position of the argument in the argument list. This feature provides for the definition  
 13102 of format strings that select arguments in an order appropriate to specific languages (see the  
 13103 EXAMPLES section).

13104 The *format* can contain either numbered argument conversion specifications (that is, "%n\$" and  
 13105 "\*m\$"), or unnumbered argument conversion specifications (that is, % and \*), but not both. The  
 13106 only exception to this is that %% can be mixed with the "%n\$" form. The results of mixing  
 13107 numbered and unnumbered argument specifications in a *format* string are undefined. When  
 13108 numbered argument specifications are used, specifying the *N*th argument requires that all the  
 13109 leading arguments, from the first to the (*N*-1)th, are specified in the format string.

13110 In format strings containing the "%n\$" form of conversion specification, numbered arguments  
 13111 in the argument list can be referenced from the format string as many times as required.

13112 In format strings containing the % form of conversion specification, each argument in the  
 13113 argument list is used exactly once.

13114 CX All forms of the *fprintf()* functions allow for the insertion of a language-dependent radix character in the output string. The radix character is defined in the program's locale (category *LC\_NUMERIC*). In the POSIX locale, or in a locale where the radix character is not defined, the radix character shall default to a period ( '.' ).

13118 XSI Each conversion specification is introduced by the '%' character or by the character sequence "%n\$", after which the following appear in sequence:

- 13120 • Zero or more *flags* (in any order), which modify the meaning of the conversion specification.
- 13121 • An optional minimum *field width*. If the converted value has fewer bytes than the field width, it shall be padded with spaces by default on the left; it shall be padded on the right if the left-adjustment flag ('-'), described below, is given to the field width. The field width takes the form of an asterisk ('\*'), described below, or a decimal integer.
- 13122
- 13123
- 13124
- 13125 • An optional *precision* that gives the minimum number of digits to appear for the d, i, o, u, x, and X conversion specifiers; the number of digits to appear after the radix character for the a, A, e, E, f, and F conversion specifiers; the maximum number of significant digits for the g and G conversion specifiers; or the maximum number of bytes to be printed from a string in s and S conversion specifiers. The precision takes the form of a period ( '.' ) followed either by an asterisk ( '\*' ), described below, or an optional decimal digit string, where a null digit string is treated as zero. If a precision appears with any other conversion specifier, the behavior is undefined.
- 13126
- 13127
- 13128 XSI
- 13129
- 13130
- 13131
- 13132
- 13133 • An optional length modifier that specifies the size of the argument.
- 13134 • A *conversion specifier* character that indicates the type of conversion to be applied.

13135 A field width, or precision, or both, may be indicated by an asterisk ( '\*' ). In this case an argument of type **int** supplies the field width or precision. Applications shall ensure that arguments specifying field width, or precision, or both appear in that order before the argument, if any, to be converted. A negative field width is taken as a '-' flag followed by a positive field width. A negative precision is taken as if the precision were omitted. In format strings containing the "%n\$" form of a conversion specification, a field width or precision may be indicated by the sequence "\*m\$", where *m* is a decimal integer in the range [1,{NL\_ARGMAX}] giving the position in the argument list (after the format argument) of an integer argument containing the field width or precision, for example:

```
13144 printf("%1$d:%2$.*3$d:%4$.*3$d\n", hour, min, precision, sec);
```

13145 The flag characters and their meanings are:

- 13146 XSI ' The integer portion of the result of a decimal conversion (%i, %d, %u, %f, %F, %g, or %G) shall be formatted with thousands' grouping characters. For other conversions the behavior is undefined. The non-monetary grouping character is used.
- 13147
- 13148
- 13149 - The result of the conversion shall be left-justified within the field. The conversion is right-justified if this flag is not specified.
- 13150
- 13151 + The result of a signed conversion shall always begin with a sign ('+' or '-'). The conversion shall begin with a sign only when a negative value is converted if this flag is not specified.
- 13152
- 13153
- 13154 <space> If the first character of a signed conversion is not a sign or if a signed conversion results in no characters, a <space> shall be prefixed to the result. This means that if the <space> and '+' flags both appear, the <space> flag shall be ignored.
- 13155
- 13156
- 13157 # Specifies that the value is to be converted to an alternative form. For o conversion, it increases the precision (if necessary) to force the first digit of the result to be zero. For x
- 13158

13159 or X conversion specifiers, a non-zero result shall have 0x (or 0X) prefixed to it. For a, A,  
 13160 e, E, f, F, g, and G conversion specifiers, the result shall always contain a radix  
 13161 character, even if no digits follow the radix character. Without this flag, a radix  
 13162 character appears in the result of these conversions only if a digit follows it. For g and G  
 13163 conversion specifiers, trailing zeros shall *not* be removed from the result as they  
 13164 normally are. For other conversion specifiers, the behavior is undefined.

13165 0 For d, i, o, u, x, X, a, A, e, E, f, F, g, and G conversion specifiers, leading zeros  
 13166 (following any indication of sign or base) are used to pad to the field width; no space  
 13167 padding is performed. If the '0' and '-' flags both appear, the '0' flag is ignored. For  
 13168 d, i, o, u, x, and X conversion specifiers, if a precision is specified, the '0' flag is  
 13169 XSI ignored. If the '0' and '\'' flags both appear, the grouping characters are inserted  
 13170 before zero padding. For other conversions, the behavior is undefined.

13171 The length modifiers and their meanings are:

13172 hh Specifies that a following d, i, o, u, x, or X conversion specifier applies to a **signed char**  
 13173 or **unsigned char** argument (the argument will have been promoted according to the  
 13174 integer promotions, but its value shall be converted to **signed char** or **unsigned char**  
 13175 before printing); or that a following n conversion specifier applies to a pointer to a  
 13176 **signed char** argument.

13177 h Specifies that a following d, i, o, u, x, or X conversion specifier applies to a **short** or  
 13178 **unsigned short** argument (the argument will have been promoted according to the  
 13179 integer promotions, but its value shall be converted to **short** or **unsigned short** before  
 13180 printing); or that a following n conversion specifier applies to a pointer to a **short**  
 13181 argument.

13182 l (ell) Specifies that a following d, i, o, u, x, or X conversion specifier applies to a **long** or  
 13183 **unsigned long** argument; that a following n conversion specifier applies to a pointer to  
 13184 a **long** argument; that a following c conversion specifier applies to a **wint\_t** argument;  
 13185 that a following s conversion specifier applies to a pointer to a **wchar\_t** argument; or  
 13186 has no effect on a following a, A, e, E, f, F, g, or G conversion specifier.

13187 ll (ell-ell) Specifies that a following d, i, o, u, x, or X conversion specifier applies to a **long long** or  
 13188 **unsigned long long** argument; or that a following n conversion specifier applies to a  
 13189 pointer to a **long long** argument.

13191 j Specifies that a following d, i, o, u, x, or X conversion specifier applies to an **intmax\_t**  
 13192 or **uintmax\_t** argument; or that a following n conversion specifier applies to a pointer  
 13193 to an **intmax\_t** argument.

13194 z Specifies that a following d, i, o, u, x, or X conversion specifier applies to a **size\_t** or the  
 13195 corresponding signed integer type argument; or that a following n conversion specifier  
 13196 applies to a pointer to a signed integer type corresponding to **size\_t** argument.

13197 t Specifies that a following d, i, o, u, x, or X conversion specifier applies to a **ptrdiff\_t** or  
 13198 the corresponding **unsigned** type argument; or that a following n conversion specifier  
 13199 applies to a pointer to a **ptrdiff\_t** argument.

13200 L Specifies that a following a, A, e, E, f, F, g, or G conversion specifier applies to a **long**  
 13201 **double** argument.

13202 If a length modifier appears with any conversion specifier other than as specified above, the  
 13203 behavior is undefined.

13204		The conversion specifiers and their meanings are:
13205	d, i	The <b>int</b> argument shall be converted to a signed decimal in the style "[−]ddd". 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 zero with an explicit precision of zero shall be no characters.
13206		
13207		
13208		
13209		
13210	o	The <b>unsigned</b> argument shall be converted to unsigned octal format in the style "ddd". 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 zero with an explicit precision of zero shall be no characters.
13211		
13212		
13213		
13214		
13215	u	The <b>unsigned</b> argument shall be converted to unsigned decimal format in the style "ddd". 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 zero with an explicit precision of zero shall be no characters.
13216		
13217		
13218		
13219		
13220	x	The <b>unsigned</b> argument shall be converted to unsigned hexadecimal format in the style "ddd"; the letters "abcdef" are used. 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 zero with an explicit precision of zero shall be no characters.
13221		
13222		
13223		
13224		
13225	X	Equivalent to the x conversion specifier, except that letters "ABCDEF" are used instead of "abcdef".
13226		
13227	f, F	The <b>double</b> argument shall be converted to decimal notation in the style "[−]ddd.ddd", where the number of digits after the radix character is equal to the precision specification. If the precision is missing, it shall be taken as 6; if the precision is explicitly zero and no '#' flag is present, no radix character shall appear. If a radix character appears, at least one digit appears before it. The low-order digit shall be rounded in an implementation-defined manner.
13228		
13229		
13230		
13231		
13232		
13233		A <b>double</b> argument representing an infinity shall be converted in one of the styles "[−]inf" or "[−]infinity"; which style is implementation-defined. A <b>double</b> argument representing a NaN shall be converted in one of the styles "[−]nan( <i>n-char-sequence</i> )"; or "[−]nan" which style, and the meaning of any <i>n-char-sequence</i> , is implementation-defined. The F conversion specifier produces "INF", "INFINITY", or "NAN" instead of "inf", "infinity", or "nan", respectively.
13234		
13235		
13236		
13237		
13238		
13239	e, E	The <b>double</b> argument shall be converted in the style "[−]d.ddde±dd", where there is one digit before the radix character (which is non-zero if the argument is non-zero) and the number of digits after it is equal to the precision; if the precision is missing, it shall be taken as 6; if the precision is zero and no '#' flag is present, no radix character shall appear. The low-order digit shall be rounded in an implementation-defined manner. The E conversion specifier shall produce a number with 'E' instead of 'e' introducing the exponent. The exponent shall always contain at least two digits. If the value is zero, the exponent shall be zero.
13240		
13241		
13242		
13243		
13244		
13245		
13246		
13247		A <b>double</b> argument representing an infinity or NaN shall be converted in the style of an f or F conversion specifier.
13248		
13249	g, G	The <b>double</b> argument shall be converted in the style f or e (or in the style E in the case of a G conversion specifier), with the precision specifying the number of significant
13250		

13251 digits. If an explicit precision is zero, it shall be taken as 1. The style used depends on |  
13252 the value converted; style `e` (or `E`) shall be used only if the exponent resulting from |  
13253 such a conversion is less than `-4` or greater than or equal to the precision. Trailing zeros |  
13254 shall be removed from the fractional portion of the result; a radix character shall appear |  
13255 only if it is followed by a digit. |

13256 A **double** argument representing an infinity or NaN shall be converted in the style of |  
13257 an `f` or `F` conversion specifier. |

13258 a, A A **double** argument representing a floating-point number shall be converted in the |  
13259 style "`[-]0xh.hhhhp±d`", where there is one hexadecimal digit (which shall be non- |  
13260 zero if the argument is a normalized floating-point number and is otherwise |  
13261 unspecified) before the decimal-point character and the number of hexadecimal digits |  
13262 after it is equal to the precision; if the precision is missing and `FLT_RADIX` is a power |  
13263 of 2, then the precision shall be sufficient for an exact representation of the value; if the |  
13264 precision is missing and `FLT_RADIX` is not a power of 2, then the precision shall be |  
13265 sufficient to distinguish values of type **double**, except that trailing zeros may be |  
13266 omitted; if the precision is zero and the `'#'` flag is not specified, no decimal-point |  
13267 character shall appear. The letters "`abcdef`" shall be used for a conversion and the |  
13268 letters "`ABCDEF`" for A conversion. The A conversion specifier produces a number with |  
13269 `'X'` and `'P'` instead of `'x'` and `'p'`. The exponent shall always contain at least one |  
13270 digit, and only as many more digits as necessary to represent the decimal exponent of |  
13271 2. If the value is zero, the exponent shall be zero. |

13272 A **double** argument representing an infinity or NaN shall be converted in the style of |  
13273 an `f` or `F` conversion specifier. |

13274 c The **int** argument shall be converted to an **unsigned char**, and the resulting byte shall |  
13275 be written. |

13276 If an `l` (`ell`) qualifier is present, the **wint\_t** argument shall be converted as if by an `l` |  
13277 conversion specification with no precision and an argument that points to a two- |  
13278 element array of type **wchar\_t**, the first element of which contains the **wint\_t** argument |  
13279 to the `l` conversion specification and the second element contains a null wide |  
13280 character. |

13281 s The argument shall be a pointer to an array of **char**. Bytes from the array shall be |  
13282 written up to (but not including) any terminating null byte. If the precision is specified, |  
13283 no more than that many bytes shall be written. If the precision is not specified or is |  
13284 greater than the size of the array, the application shall ensure that the array contains a |  
13285 null byte. |

13286 If an `l` (`ell`) qualifier is present, the argument shall be a pointer to an array of type |  
13287 **wchar\_t**. Wide characters from the array shall be converted to characters (each as if by |  
13288 a call to the `wcrtomb()` function, with the conversion state described by an **mbstate\_t** |  
13289 object initialized to zero before the first wide character is converted) up to and |  
13290 including a terminating null wide character. The resulting characters shall be written |  
13291 up to (but not including) the terminating null character (byte). If no precision is |  
13292 specified, the application shall ensure that the array contains a null wide character. If a |  
13293 precision is specified, no more than that many characters (bytes) shall be written |  
13294 (including shift sequences, if any), and the array shall contain a null wide character if, |  
13295 to equal the character sequence length given by the precision, the function would need |  
13296 to access a wide character one past the end of the array. In no case shall a partial |  
13297 character written. |

13298	p	The argument shall be a pointer to <b>void</b> . The value of the pointer is converted to a sequence of printable characters, in an implementation-defined manner.	
13299			
13300	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.	
13301			
13302			
13303	XSI C	Equivalent to <code>lc</code> .	
13304	XSI S	Equivalent to <code>ls</code> .	
13305	%	Print a '%' character; no argument is converted. The complete conversion specification shall be <code>%%</code> .	
13306			
13307		If a conversion specification does not match one of the above forms, the behavior is undefined. If any argument is not the correct type for the corresponding conversion specification, the behavior is undefined.	
13308			
13309			
13310		In no case shall a nonexistent or small field width cause truncation of a field; if the result of a conversion is wider than the field width, the field shall be expanded to contain the conversion result. Characters generated by <i>fprintf()</i> and <i>printf()</i> are printed as if <i>fputc()</i> had been called.	
13311			
13312			
13313		For the <code>a</code> and <code>A</code> conversion specifiers, if <code>FLT_RADIX</code> is a power of 2, the value shall be correctly rounded to a hexadecimal floating number with the given precision.	
13314			
13315		For <code>a</code> and <code>A</code> conversions, if <code>FLT_RADIX</code> is not a power of 2 and the result is not exactly representable in the given precision, the result should be one of the two adjacent numbers in hexadecimal floating style with the given precision, with the extra stipulation that the error should have a correct sign for the current rounding direction.	
13316			
13317			
13318			
13319		For the <code>e</code> , <code>E</code> , <code>f</code> , <code>F</code> , <code>g</code> , and <code>G</code> conversion specifiers, if the number of significant decimal digits is at most <code>DECIMAL_DIG</code> , then the result should be correctly rounded. If the number of significant decimal digits is more than <code>DECIMAL_DIG</code> but the source value is exactly representable with <code>DECIMAL_DIG</code> digits, then the result should be an exact representation with trailing zeros. Otherwise, the source value is bounded by two adjacent decimal strings $L < U$ , both having <code>DECIMAL_DIG</code> significant digits; the value of the resultant decimal string $D$ should satisfy $L \leq D \leq U$ , with the extra stipulation that the error should have a correct sign for the current rounding direction.	
13320			
13321			
13322			
13323			
13324			
13325			
13326			
13327	CX	The <code>st_ctime</code> and <code>st_mtime</code> fields of the file shall be marked for update between the call to a successful execution of <i>fprintf()</i> or <i>printf()</i> and the next successful completion of a call to <i>fflush()</i> or <i>fclose()</i> on the same stream or a call to <i>exit()</i> or <i>abort()</i> .	
13328			
13329			
13330	<b>RETURN VALUE</b>		
13331		Upon successful completion, the <i>fprintf()</i> and <i>printf()</i> functions shall return the number of bytes transmitted.	
13332			
13333		Upon successful completion, the <i>sprintf()</i> function shall return the number of bytes written to <code>s</code> , excluding the terminating null byte.	
13334			
13335		Upon successful completion, the <i>snprintf()</i> function shall return the number of bytes that would be written to <code>s</code> had <code>n</code> been sufficiently large excluding the terminating null byte.	
13336			
13337		If an output error was encountered, these functions shall return a negative value.	
13338		If the value of <code>n</code> is zero on a call to <i>snprintf()</i> , nothing shall be written, the number of bytes that would have been written had <code>n</code> been sufficiently large excluding the terminating null shall be returned, and <code>s</code> may be a null pointer.	
13339			
13340			



13341 **ERRORS**

13342 For the conditions under which *fprintf()* and *printf()* fail and may fail, refer to *fputc()* or  
13343 *fputwc()*.

13344 In addition, all forms of *fprintf()* may fail if:

13345 XSI [EILSEQ] A wide-character code that does not correspond to a valid character has been  
13346 detected.

13347 XSI [EINVAL] There are insufficient arguments.

13348 The *printf()* and *fprintf()* functions may fail if:

13349 XSI [ENOMEM] Insufficient storage space is available.

13350 The *snprintf()* function shall fail if:

13351 XSI [EOVERFLOW] The value of *n* is greater than {INT\_MAX} or the number of bytes needed to  
13352 hold the output excluding the terminating null is greater than {INT\_MAX}.

13353 **EXAMPLES**13354 **Printing Language-Independent Date and Time**

13355 The following statement can be used to print date and time using a language-independent  
13356 format:

```
13357 printf(format, weekday, month, day, hour, min);
```

13358 For American usage, *format* could be a pointer to the following string:

```
13359 "%s, %s %d, %d:%.2d\n"
```

13360 This example would produce the following message:

```
13361 Sunday, July 3, 10:02
```

13362 For German usage, *format* could be a pointer to the following string:

```
13363 "%1$s, %3$d. %2$s, %4$d:%5$.2d\n"
```

13364 This definition of *format* would produce the following message:

```
13365 Sonntag, 3. Juli, 10:02
```

13366 **Printing File Information**

13367 The following example prints information about the type, permissions, and number of links of a  
13368 specific file in a directory.

13369 The first two calls to *printf()* use data decoded from a previous *stat()* call. The user-defined  
13370 *strperm()* function shall return a string similar to the one at the beginning of the output for the  
13371 following command:

```
13372 ls -l
```

13373 The next call to *printf()* outputs the owner's name if it is found using *getpwuid()*; the *getpwuid()*  
13374 function shall return a **passwd** structure from which the name of the user is extracted. If the user  
13375 name is not found, the program instead prints out the numeric value of the user ID.

13376 The next call prints out the group name if it is found using *getgrgid()*; *getgrgid()* is very similar to  
13377 *getpwuid()* except that it shall return group information based on the group number. Once  
13378 again, if the group is not found, the program prints the numeric value of the group for the entry.

```

13379     The final call to printf() prints the size of the file.
13380     #include <stdio.h>
13381     #include <sys/types.h>
13382     #include <pwd.h>
13383     #include <grp.h>
13384     char *strperm (mode_t);
13385     ...
13386     struct stat statbuf;
13387     struct passwd *pwd;
13388     struct group *grp;
13389     ...
13390     printf("%10.10s", strperm (statbuf.st_mode));
13391     printf("%4d", statbuf.st_nlink);
13392     if ((pwd = getpwuid(statbuf.st_uid)) != NULL)
13393         printf(" %-8.8s", pwd->pw_name);
13394     else
13395         printf(" %-8ld", (long) statbuf.st_uid);
13396     if ((grp = getgrgid(statbuf.st_gid)) != NULL)
13397         printf(" %-8.8s", grp->gr_name);
13398     else
13399         printf(" %-8ld", (long) statbuf.st_gid);
13400     printf("%9jd", (intmax_t) statbuf.st_size);
13401     ...

```

#### 13402 **Printing a Localized Date String**

13403 The following example gets a localized date string. The *nl\_langinfo()* function shall return the  
 13404 localized date string, which specifies the order and layout of the date. The *strftime()* function  
 13405 takes this information and, using the **tm** structure for values, places the date and time  
 13406 information into *datestring*. The *printf()* function then outputs *datestring* and the name of the  
 13407 entry.

```

13408     #include <stdio.h>
13409     #include <time.h>
13410     #include <langinfo.h>
13411     ...
13412     struct dirent *dp;
13413     struct tm *tm;
13414     char datestring[256];
13415     ...
13416     strftime(datestring, sizeof(datestring), nl_langinfo (D_T_FMT), tm);
13417     printf(" %s %s\n", datestring, dp->d_name);
13418     ...

```

13419 **Printing Error Information**

13420 The following example uses *fprintf()* to write error information to standard error.

13421 In the first group of calls, the program tries to open the password lock file named **LOCKFILE**. If  
 13422 the file already exists, this is an error, as indicated by the **O\_EXCL** flag on the *open()* function. If  
 13423 the call fails, the program assumes that someone else is updating the password file, and the  
 13424 program exits.

13425 The next group of calls saves a new password file as the current password file by creating a link  
 13426 between **LOCKFILE** and the new password file **PASSWDFILE**.

```

13427 #include <sys/types.h>
13428 #include <sys/stat.h>
13429 #include <fcntl.h>
13430 #include <stdio.h>
13431 #include <stdlib.h>
13432 #include <unistd.h>
13433 #include <string.h>
13434 #include <errno.h>

13435 #define LOCKFILE "/etc/ptmp"
13436 #define PASSWDFILE "/etc/passwd"
13437 ...
13438 int pfd;
13439 ...
13440 if ((pfd = open(LOCKFILE, O_WRONLY | O_CREAT | O_EXCL,
13441               S_IRUSR | S_IWUSR | S_IRGRP | S_IROTH)) == -1)
13442 {
13443     fprintf(stderr, "Cannot open /etc/ptmp. Try again later.\n");
13444     exit(1);
13445 }
13446 ...
13447 if (link(LOCKFILE, PASSWDFILE) == -1) {
13448     fprintf(stderr, "Link error: %s\n", strerror(errno));
13449     exit(1);
13450 }
13451 ...

```

13452 **Printing Usage Information**

13453 The following example checks to make sure the program has the necessary arguments, and uses  
 13454 *fprintf()* to print usage information if the expected number of arguments is not present.

```

13455 #include <stdio.h>
13456 #include <stdlib.h>
13457 ...
13458 char *Options = "hdbt1";
13459 ...
13460 if (argc < 2) {
13461     fprintf(stderr, "Usage: %s -%s <file\n", argv[0], Options); exit(1);
13462 }
13463 ...

```

13464       **Formatting a Decimal String**

13465       The following example prints a key and data pair on *stdout*. Note use of the '\*' (asterisk) in the  
13466       format string; this ensures the correct number of decimal places for the element based on the  
13467       number of elements requested.

```
13468       #include <stdio.h>
13469       ...
13470       long i;
13471       char *keyst;
13472       int elementlen, len;
13473       ...
13474       while (len < elementlen) {
13475       ...
13476           printf("%s Element%0*ld\n", keyst, elementlen, i);
13477       ...
13478       }
```

13479       **Creating a Filename**

13480       The following example creates a filename using information from a previous *getpwnam()*  
13481       function that returned the HOME directory of the user.

```
13482       #include <stdio.h>
13483       #include <sys/types.h>
13484       #include <unistd.h>
13485       ...
13486       char filename[PATH_MAX+1];
13487       struct passwd *pw;
13488       ...
13489       sprintf(filename, "%s/%d.out", pw->pw_dir, getpid());
13490       ...
```

13491       **Reporting an Event**

13492       The following example loops until an event has timed out. The *pause()* function waits forever  
13493       unless it receives a signal. The *fprintf()* statement should never occur due to the possible return  
13494       values of *pause()*.

```
13495       #include <stdio.h>
13496       #include <unistd.h>
13497       #include <string.h>
13498       #include <errno.h>
13499       ...
13500       while (!event_complete) {
13501       ...
13502           if (pause() != -1 || errno != EINTR)
13503               fprintf(stderr, "pause: unknown error: %s\n", strerror(errno));
13504       }
13505       ...
```

13506 **Printing Monetary Information**

13507 The following example uses *strfmon()* to convert a number and store it as a formatted monetary  
 13508 string named *convbuf*. If the first number is printed, the program prints the format and the  
 13509 description; otherwise, it just prints the number.

```

13510 #include <monetary.h>
13511 #include <stdio.h>
13512 ...
13513 struct tblfmt {
13514     char *format;
13515     char *description;
13516 };
13517 struct tblfmt table[] = {
13518     { "%n", "default formatting" },
13519     { "%11n", "right align within an 11 character field" },
13520     { "%#5n", "aligned columns for values up to 99,999" },
13521     { "%=*#5n", "specify a fill character" },
13522     { "%=0#5n", "fill characters do not use grouping" },
13523     { "%^#5n", "disable the grouping separator" },
13524     { "%^#5.0n", "round off to whole units" },
13525     { "%^#5.4n", "increase the precision" },
13526     { "%(#5n", "use an alternative pos/neg style" },
13527     { "%!(#5n", "disable the currency symbol" },
13528 };
13529 ...
13530 float input[3];
13531 int i, j;
13532 char convbuf[100];
13533 ...
13534 strfmon(convbuf, sizeof(convbuf), table[i].format, input[j]);
13535
13536 if (j == 0) {
13537     printf("%s%s%s\n", table[i].format,
13538         convbuf, table[i].description);
13539 }
13540 else {
13541     printf("%s\n", convbuf);
13542 }
13543 ...

```

13543 **APPLICATION USAGE**

13544 If the application calling *fprintf()* has any objects of type **wint\_t** or **wchar\_t**, it must also include  
 13545 the **<wchar.h>** header to have these objects defined.

13546 **RATIONALE**

13547 None.

13548 **FUTURE DIRECTIONS**

13549 None.

13550 **SEE ALSO**

13551 *fputc()*, *fscanf()*, *setlocale()*, *wcrtomb()*, the Base Definitions volume of IEEE Std 1003.1-200x,  
 13552 **<stdio.h>**, **<wchar.h>**, the Base Definitions volume of IEEE Std 1003.1-200x, Chapter 7, Locale

13553 **CHANGE HISTORY**

13554 First released in Issue 1. Derived from Issue 1 of the SVID.

13555 **Issue 5**

13556 Aligned with ISO/IEC 9899:1990/Amendment 1:1995 (E). Specifically, the l (ell) qualifier can  
13557 now be used with c and s conversion specifiers.

13558 The *snprintf()* function is new in Issue 5.

13559 **Issue 6**

13560 Extensions beyond the ISO C standard are now marked.

13561 The DESCRIPTION is updated to avoid use of the term “must” for application requirements. |

13562 The following changes are made for alignment with the ISO/IEC 9899:1999 standard:

13563 • The prototypes for *fprintf()*, *printf()*, *snprintf()*, and *sprintf()* are updated, and the XSI  
13564 shading is removed from *snprintf()*.

13565 • The description of *snprintf()* is aligned with the ISO C standard. Note that this supersedes |  
13566 the *snprintf()* description in The Open Group Base Resolution bwg98-006, which changed the |  
13567 behavior from Issue 5. |

13568 • The DESCRIPTION is updated. |

13569 The DESCRIPTION is updated to use the terms “conversion specifier” and “conversion  
13570 specification” consistently. |

13571 ISO/IEC 9899:1999 standard, Technical Corrigendum No. 1 is incorporated. |

13572 **NAME**

13573       fputc — put a byte on a stream

13574 **SYNOPSIS**

13575       #include &lt;stdio.h&gt;

13576       int fputc(int *c*, FILE \**stream*);13577 **DESCRIPTION**

13578 CX       The functionality described on this reference page is aligned with the ISO C standard. Any  
 13579       conflict between the requirements described here and the ISO C standard is unintentional. This  
 13580       volume of IEEE Std 1003.1-200x defers to the ISO C standard.

13581       The *fputc()* function shall write the byte specified by *c* (converted to an **unsigned char**) to the  
 13582       output stream pointed to by *stream*, at the position indicated by the associated file-position  
 13583       indicator for the stream (if defined), and shall advance the indicator appropriately. If the file  
 13584       cannot support positioning requests, or if the stream was opened with append mode, the byte  
 13585       shall be appended to the output stream.

13586 CX       The *st\_ctime* and *st\_mtime* fields of the file shall be marked for update between the successful  
 13587       execution of *fputc()* and the next successful completion of a call to *fflush()* or *fclose()* on the same  
 13588       stream or a call to *exit()* or *abort()*.

13589 **RETURN VALUE**

13590       Upon successful completion, *fputc()* shall return the value it has written. Otherwise, it shall  
 13591 CX       return EOF, the error indicator for the stream shall be set, and *errno* shall be set to indicate the  
 13592       error.

13593 **ERRORS**

13594       The *fputc()* function shall fail if either the *stream* is unbuffered or the *stream*'s buffer needs to be  
 13595       flushed, and:

13596 CX       [EAGAIN]       The O\_NONBLOCK flag is set for the file descriptor underlying *stream* and the  
 13597       process would be delayed in the write operation.

13598 CX       [EBADF]       The file descriptor underlying *stream* is not a valid file descriptor open for  
 13599       writing.

13600 CX       [EFBIG]       An attempt was made to write to a file that exceeds the maximum file size.

13601 XSI       [EFBIG]       An attempt was made to write to a file that exceeds the process' file size limit.

13602 CX       [EFBIG]       The file is a regular file and an attempt was made to write at or beyond the  
 13603       offset maximum.

13604 CX       [EINTR]       The write operation was terminated due to the receipt of a signal, and no data  
 13605       was transferred.

13606 CX       [EIO]        A physical I/O error has occurred, or the process is a member of a  
 13607       background process group attempting to write to its controlling terminal,  
 13608       TOSTOP is set, the process is neither ignoring nor blocking SIGTTOU, and the  
 13609       process group of the process is orphaned. This error may also be returned  
 13610       under implementation-defined conditions.

13611 CX       [ENOSPC]       There was no free space remaining on the device containing the file.

13612 CX       [EPIPE]       An attempt is made to write to a pipe or FIFO that is not open for reading by  
 13613       any process. A SIGPIPE signal shall also be sent to the thread.

13614       The *fputc()* function may fail if:

13615 CX [ENOMEM] Insufficient storage space is available.

13616 CX [ENXIO] A request was made of a nonexistent device, or the request was outside the  
13617 capabilities of the device.

13618 **EXAMPLES**

13619 None.

13620 **APPLICATION USAGE**

13621 None.

13622 **RATIONALE**

13623 None.

13624 **FUTURE DIRECTIONS**

13625 None.

13626 **SEE ALSO**

13627 *ferror()*, *fopen()*, *getrlimit()*, *putc()*, *puts()*, *setbuf()*, *ulimit()*, the Base Definitions volume of  
13628 IEEE Std 1003.1-200x, <stdio.h>

13629 **CHANGE HISTORY**

13630 First released in Issue 1. Derived from Issue 1 of the SVID.

13631 **Issue 5**

13632 Large File Summit extensions are added.

13633 **Issue 6**

13634 Extensions beyond the ISO C standard are now marked.

13635 The following new requirements on POSIX implementations derive from alignment with the  
13636 Single UNIX Specification:

- 13637
- The [EIO] and [EFBIG] mandatory error conditions are added.
- 13638
- The [ENOMEM] and [ENXIO] optional error conditions are added.



13639 **NAME**

13640 fputs — put a string on a stream

13641 **SYNOPSIS**

13642 #include &lt;stdio.h&gt;

13643 int fputs(const char \*restrict *s*, FILE \*restrict *stream*);13644 **DESCRIPTION**

13645 cx The functionality described on this reference page is aligned with the ISO C standard. Any  
 13646 conflict between the requirements described here and the ISO C standard is unintentional. This  
 13647 volume of IEEE Std 1003.1-200x defers to the ISO C standard.

13648 The *fputs()* function shall write the null-terminated string pointed to by *s* to the stream pointed  
 13649 to by *stream*. The terminating null byte shall not be written.

13650 cx The *st\_ctime* and *st\_mtime* fields of the file shall be marked for update between the successful  
 13651 execution of *fputs()* and the next successful completion of a call to *fflush()* or *fclose()* on the same  
 13652 stream or a call to *exit()* or *abort()*.

13653 **RETURN VALUE**

13654 Upon successful completion, *fputs()* shall return a non-negative number. Otherwise, it shall  
 13655 cx return EOF, set an error indicator for the stream, and set *errno* to indicate the error.

13656 **ERRORS**13657 Refer to *fputc()*.13658 **EXAMPLES**13659 **Printing to Standard Output**

13660 The following example gets the current time, converts it to a string using *localtime()* and  
 13661 *asctime()*, and prints it to standard output using *fputs()*. It then prints the number of minutes to  
 13662 an event for which it is waiting.

```

13663 #include <time.h>
13664 #include <stdio.h>
13665 ...
13666 time_t now;
13667 int minutes_to_event;
13668 ...
13669 time(&now);
13670 printf("The time is ");
13671 fputs(asctime(localtime(&now)), stdout);
13672 printf("There are still %d minutes to the event.\n",
13673        minutes_to_event);
13674 ...

```

13675 **APPLICATION USAGE**13676 The *puts()* function appends a <newline> while *fputs()* does not.13677 **RATIONALE**

13678 None.

13679 **FUTURE DIRECTIONS**

13680 None.

13681 **SEE ALSO**

13682        *fopen()*, *putc()*, *puts()*, the Base Definitions volume of IEEE Std 1003.1-200x, <stdio.h>

13683 **CHANGE HISTORY**

13684        First released in Issue 1. Derived from Issue 1 of the SVID.

13685 **Issue 6**

13686        Extensions beyond the ISO C standard are now marked.

13687        The *fputs()* prototype is updated for alignment with the ISO/IEC 9899: 1999 standard.

13688 **NAME**13689 `fputc` — put a wide-character code on a stream13690 **SYNOPSIS**13691 `#include <stdio.h>`13692 `#include <wchar.h>`13693 `wint_t fputc(wchar_t wc, FILE *stream);`13694 **DESCRIPTION**

13695 CX The functionality described on this reference page is aligned with the ISO C standard. Any  
 13696 conflict between the requirements described here and the ISO C standard is unintentional. This  
 13697 volume of IEEE Std 1003.1-200x defers to the ISO C standard.

13698 The `fputc()` function shall write the character corresponding to the wide-character code `wc` to  
 13699 the output stream pointed to by `stream`, at the position indicated by the associated file-position  
 13700 indicator for the stream (if defined), and advances the indicator appropriately. If the file cannot  
 13701 support positioning requests, or if the stream was opened with append mode, the character is  
 13702 appended to the output stream. If an error occurs while writing the character, the shift state of  
 13703 the output file is left in an undefined state.

13704 CX The `st_ctime` and `st_mtime` fields of the file shall be marked for update between the successful  
 13705 execution of `fputc()` and the next successful completion of a call to `fflush()` or `fclose()` on the  
 13706 same stream or a call to `exit()` or `abort()`.

13707 **RETURN VALUE**

13708 Upon successful completion, `fputc()` shall return `wc`. Otherwise, it shall return WEOF, the error  
 13709 CX indicator for the stream shall be set set, and `errno` shall be set to indicate the error.

13710 **ERRORS**

13711 The `fputc()` function shall fail if either the stream is unbuffered or data in the `stream`'s buffer  
 13712 needs to be written, and:

13713 CX [EAGAIN] The O\_NONBLOCK flag is set for the file descriptor underlying `stream` and the  
 13714 process would be delayed in the write operation.

13715 CX [EBADF] The file descriptor underlying `stream` is not a valid file descriptor open for  
 13716 writing.

13717 CX [EFBIG] An attempt was made to write to a file that exceeds the maximum file size or  
 13718 the process' file size limit.

13719 CX [EFBIG] The file is a regular file and an attempt was made to write at or beyond the  
 13720 offset maximum associated with the corresponding stream.

13721 [EILSEQ] The wide-character code `wc` does not correspond to a valid character.

13722 CX [EINTR] The write operation was terminated due to the receipt of a signal, and no data  
 13723 was transferred.

13724 CX [EIO] A physical I/O error has occurred, or the process is a member of a  
 13725 background process group attempting to write to its controlling terminal,  
 13726 TOSTOP is set, the process is neither ignoring nor blocking SIGTTOU, and the  
 13727 process group of the process is orphaned. This error may also be returned  
 13728 under implementation-defined conditions.

13729 CX [ENOSPC] There was no free space remaining on the device containing the file.

13730 CX [EPIPE] An attempt is made to write to a pipe or FIFO that is not open for reading by  
 13731 any process. A SIGPIPE signal shall also be sent to the thread.

13732 The *fputc()* function may fail if:

13733 CX [ENOMEM] Insufficient storage space is available.

13734 CX [ENXIO] A request was made of a nonexistent device, or the request was outside the  
13735 capabilities of the device.

#### 13736 EXAMPLES

13737 None.

#### 13738 APPLICATION USAGE

13739 None.

#### 13740 RATIONALE

13741 None.

#### 13742 FUTURE DIRECTIONS

13743 None.

#### 13744 SEE ALSO

13745 *ferror()*, *fopen()*, *setbuf()*, *ulimit()*, the Base Definitions volume of IEEE Std 1003.1-200x,  
13746 `<stdio.h>`, `<wchar.h>`

#### 13747 CHANGE HISTORY

13748 First released in Issue 4. Derived from the MSE working draft.

#### 13749 Issue 5

13750 Aligned with ISO/IEC 9899:1990/Amendment 1:1995 (E). Specifically, the type of argument *wc*  
13751 is changed from `wint_t` to `wchar_t`.

13752 The Optional Header (OH) marking is removed from `<stdio.h>`.

13753 Large File Summit extensions are added.

#### 13754 Issue 6

13755 Extensions beyond the ISO C standard are now marked.

13756 The following new requirements on POSIX implementations derive from alignment with the  
13757 Single UNIX Specification:

- 13758 • The [EFBIG] and [EIO] mandatory error conditions are added.
- 13759 • The [ENOMEM] and [ENXIO] optional error conditions are added.

13760 **NAME**

13761 fputws — put a wide-character string on a stream

13762 **SYNOPSIS**

13763 #include &lt;stdio.h&gt;

13764 #include &lt;wchar.h&gt;

13765 int fputws(const wchar\_t \*restrict ws, FILE \*restrict stream);

13766 **DESCRIPTION**

13767 CX The functionality described on this reference page is aligned with the ISO C standard. Any  
 13768 conflict between the requirements described here and the ISO C standard is unintentional. This  
 13769 volume of IEEE Std 1003.1-200x defers to the ISO C standard.

13770 The *fputws()* function shall write a character string corresponding to the (null-terminated)  
 13771 wide-character string pointed to by *ws* to the stream pointed to by *stream*. No character  
 13772 corresponding to the terminating null wide-character code shall be written.

13773 CX The *st\_ctime* and *st\_mtime* fields of the file shall be marked for update between the successful  
 13774 execution of *fputws()* and the next successful completion of a call to *flush()* or *fclose()* on the  
 13775 same stream or a call to *exit()* or *abort()*.

13776 **RETURN VALUE**

13777 Upon successful completion, *fputws()* shall return a non-negative number. Otherwise, it shall  
 13778 CX return  $-1$ , set an error indicator for the stream, and set *errno* to indicate the error.

13779 **ERRORS**13780 Refer to *fputwc()*.13781 **EXAMPLES**

13782 None.

13783 **APPLICATION USAGE**13784 The *fputws()* function does not append a <newline>.13785 **RATIONALE**

13786 None.

13787 **FUTURE DIRECTIONS**

13788 None.

13789 **SEE ALSO**13790 *fopen()*, the Base Definitions volume of IEEE Std 1003.1-200x, <stdio.h>, <wchar.h>13791 **CHANGE HISTORY**

13792 First released in Issue 4. Derived from the MSE working draft.

13793 **Issue 5**

13794 The Optional Header (OH) marking is removed from &lt;stdio.h&gt;.

13795 **Issue 6**

13796 Extensions beyond the ISO C standard are now marked.

13797 The *fputws()* prototype is updated for alignment with the ISO/IEC 9899:1999 standard.

## 13798 NAME

13799 fread — binary input

## 13800 SYNOPSIS

13801 #include &lt;stdio.h&gt;

```
13802 size_t fread(void *restrict ptr, size_t size, size_t nitems,
13803 FILE *restrict stream);
```

## 13804 DESCRIPTION

13805 CX The functionality described on this reference page is aligned with the ISO C standard. Any  
 13806 conflict between the requirements described here and the ISO C standard is unintentional. This  
 13807 volume of IEEE Std 1003.1-200x defers to the ISO C standard.

13808 The *fread()* function shall read into the array pointed to by *ptr* up to *nitems* elements whose size  
 13809 is specified by *size* in bytes, from the stream pointed to by *stream*. For each object, *size* calls shall  
 13810 be made to the *fgetc()* function and the results stored, in the order read, in an array of **unsigned**  
 13811 **char** exactly overlaying the object. The file position indicator for the stream (if defined) shall be  
 13812 advanced by the number of bytes successfully read. If an error occurs, the resulting value of the  
 13813 file position indicator for the stream is unspecified. If a partial element is read, its value is  
 13814 unspecified.

13815 CX The *fread()* function may mark the *st\_atime* field of the file associated with *stream* for update. The  
 13816 *st\_atime* field shall be marked for update by the first successful execution of *fgetc()*, *fgets()*,  
 13817 *fgetwc()*, *fgetws()*, *fread()*, *fscanf()*, *getc()*, *getchar()*, *gets()*, or *scanf()* using *stream* that returns  
 13818 data not supplied by a prior call to *ungetc()* or *ungetwc()*.

## 13819 RETURN VALUE

13820 Upon successful completion, *fread()* shall return the number of elements successfully read which  
 13821 is less than *nitems* only if a read error or end-of-file is encountered. If *size* or *nitems* is 0, *fread()*  
 13822 shall return 0 and the contents of the array and the state of the stream remain unchanged.  
 13823 CX Otherwise, if a read error occurs, the error indicator for the stream shall be set, and *errno* shall be  
 13824 set to indicate the error.

## 13825 ERRORS

13826 Refer to *fgetc()*.

## 13827 EXAMPLES

13828 **Reading from a Stream**13829 The following example reads a single element from the *fp* stream into the array pointed to by *buf*.

```
13830 #include <stdio.h>
13831 ...
13832 size_t bytes_read;
13833 char buf[100];
13834 FILE *fp;
13835 ...
13836 bytes_read = fread(buf, sizeof(buf), 1, fp);
13837 ...
```

## 13838 APPLICATION USAGE

13839 The *ferror()* or *feof()* functions must be used to distinguish between an error condition and an  
 13840 end-of-file condition.

13841 Because of possible differences in element length and byte ordering, files written using *fwrite()*  
 13842 are application-dependent, and possibly cannot be read using *fread()* by a different application

13843 or by the same application on a different processor.

13844 **RATIONALE**

13845 None.

13846 **FUTURE DIRECTIONS**

13847 None.

13848 **SEE ALSO**

13849 *feof()*, *ferror()*, *fgetc()*, *fopen()*, *getc()*, *gets()*, *scanf()*, the Base Definitions volume of  
13850 IEEE Std 1003.1-200x, <**stdio.h**>

13851 **CHANGE HISTORY**

13852 First released in Issue 1. Derived from Issue 1 of the SVID.

13853 **Issue 6**

13854 Extensions beyond the ISO C standard are now marked.

13855 The following changes are made for alignment with the ISO/IEC 9899:1999 standard:

- 13856
- The *fread()* prototype is updated.
- 13857
- The DESCRIPTION is updated to describe how the bytes from a call to *fgetc()* are stored.

13858 **NAME**

13859 free — free allocated memory

13860 **SYNOPSIS**

13861 #include &lt;stdlib.h&gt;

13862 void free(void \*ptr);

13863 **DESCRIPTION**

13864 CX The functionality described on this reference page is aligned with the ISO C standard. Any  
13865 conflict between the requirements described here and the ISO C standard is unintentional. This  
13866 volume of IEEE Std 1003.1-200x defers to the ISO C standard.

13867 The *free()* function shall cause the space pointed to by *ptr* to be deallocated; that is, made  
13868 available for further allocation. If *ptr* is a null pointer, no action shall occur. Otherwise, if the  
13869 ADV argument does not match a pointer earlier returned by the *calloc()*, *malloc()*, *posix\_memalign()*,  
13870 XSI *realloc()*, or *strdup()*, function, or if the space has been deallocated by a call to *free()* or *realloc()*,  
13871 the behavior is undefined.

13872 Any use of a pointer that refers to freed space results in undefined behavior.

13873 **RETURN VALUE**13874 The *free()* function shall not return a value.13875 **ERRORS**

13876 No errors are defined.

13877 **EXAMPLES**

13878 None.

13879 **APPLICATION USAGE**

13880 There is now no requirement for the implementation to support the inclusion of &lt;malloc.h&gt;.

13881 **RATIONALE**

13882 None.

13883 **FUTURE DIRECTIONS**

13884 None.

13885 **SEE ALSO**13886 *calloc()*, *malloc()*, *realloc()*, the Base Definitions volume of IEEE Std 1003.1-200x, <stdlib.h>13887 **CHANGE HISTORY**

13888 First released in Issue 1. Derived from Issue 1 of the SVID.

13889 **Issue 6**13890 Reference to the *valloc()* function is removed.



## 13891 NAME

13892 freeaddrinfo, getaddrinfo — get address information

## 13893 SYNOPSIS

```

13894 #include <sys/socket.h>
13895 #include <netdb.h>

13896 void freeaddrinfo(struct addrinfo *ai);
13897 int getaddrinfo(const char *restrict nodename,
13898                const char *restrict servname,
13899                const struct addrinfo *restrict hints,
13900                struct addrinfo **restrict res);

```

## 13901 DESCRIPTION

13902 The *freeaddrinfo()* function shall free one or more **addrinfo** structures returned by *getaddrinfo()*, |  
 13903 along with any additional storage associated with those structures. If the *ai\_next* field of the |  
 13904 structure is not null, the entire list of structures shall be freed. The *freeaddrinfo()* function shall |  
 13905 support the freeing of arbitrary sublists of an **addrinfo** list originally returned by *getaddrinfo()*.

13906 The *getaddrinfo()* function shall translate the name of a service location (for example, a host |  
 13907 name) and/or a service name and shall return a set of socket addresses and associated |  
 13908 information to be used in creating a socket with which to address the specified service.

13909 The *freeaddrinfo()* and *getaddrinfo()* functions shall be thread-safe.

13910 The *nodename* and *servname* arguments are either null pointers or pointers to null-terminated |  
 13911 strings. One or both of these two arguments shall be supplied by the application as a non-null |  
 13912 pointer.

13913 The format of a valid name depends on the address family or families. If a specific family is not |  
 13914 given and the name could be interpreted as valid within multiple supported families, the |  
 13915 implementation shall attempt to resolve the name in all supported families and, in absence of |  
 13916 errors, one or more results shall be returned.

13917 If the *nodename* argument is not null, it can be a descriptive name or can be an address string. If |  
 13918 IP6 the specified address family is AF\_INET, AF\_INET6, or AF\_UNSPEC, valid descriptive names |  
 13919 include host names. If the specified address family is AF\_INET or AF\_UNSPEC, address strings |  
 13920 using Internet standard dot notation as specified in *inet\_addr()* are valid.

13921 IP6 If the specified address family is AF\_INET6 or AF\_UNSPEC, standard IPv6 text forms described |  
 13922 in *inet\_ntop()* are valid.

13923 If *nodename* is not null, the requested service location is named by *nodename*; otherwise, the |  
 13924 requested service location is local to the caller.

13925 If *servname* is null, the call shall return network-level addresses for the specified *nodename*. If |  
 13926 *servname* is not null, it is a null-terminated character string identifying the requested service. This |  
 13927 can be either a descriptive name or a numeric representation suitable for use with the address |  
 13928 IP6 family or families. If the specified address family is AF\_INET, AF\_INET6, or AF\_UNSPEC, the |  
 13929 service can be specified as a string specifying a decimal port number.

13930 If the *hints* argument is not null, it refers to a structure containing input values that may direct |  
 13931 the operation by providing options and by limiting the returned information to a specific socket |  
 13932 type, address family and/or protocol. In this *hints* structure every member other than *ai\_flags*, |  
 13933 *ai\_family*, *ai\_socktype*, and *ai\_protocol* shall be set to zero or a null pointer. A value of |  
 13934 AF\_UNSPEC for *ai\_family* means that the caller shall accept any address family. A value of zero |  
 13935 for *ai\_socktype* means that the caller shall accept any socket type. A value of zero for *ai\_protocol* |  
 13936 means that the caller shall accept any protocol. If *hints* is a null pointer, the behavior shall be as if

13937 it referred to a structure containing the value zero for the *ai\_flags*, *ai\_socktype*, and *ai\_protocol*  
13938 fields, and AF\_UNSPEC for the *ai\_family* field.

13939 The *ai\_flags* field to which the *hints* parameter points shall be set to zero or be the bitwise-  
13940 inclusive OR of one or more of the values AI\_PASSIVE, AI\_CANONNAME,  
13941 AI\_NUMERICHOST, and AI\_NUMERICSERV.

13942 If the AI\_PASSIVE flag is specified, the returned address information shall be suitable for use in  
13943 binding a socket for accepting incoming connections for the specified service. In this case, if the  
13944 *nodename* argument is null, then the IP address portion of the socket address structure shall be  
13945 set to INADDR\_ANY for an IPv4 address or IN6ADDR\_ANY\_INIT for an IPv6 address. If the  
13946 AI\_PASSIVE flag is not specified, the returned address information shall be suitable for a call to  
13947 *connect()* (for a connection-mode protocol) or for a call to *connect()*, *sendto()*, or *sendmsg()* (for a  
13948 connectionless protocol). In this case, if the *nodename* argument is null, then the IP address  
13949 portion of the socket address structure shall be set to the loopback address.

13950 If the AI\_CANONNAME flag is specified and the *nodename* argument is not null, the function  
13951 shall attempt to determine the canonical name corresponding to *nodename* (for example, if  
13952 *nodename* is an alias or shorthand notation for a complete name).

13953 If the AI\_NUMERICHOST flag is specified, then a non-null *nodename* string supplied shall be a  
13954 numeric host address string. Otherwise, an [EAI\_NONAME] error is returned. This flag shall  
13955 prevent any type of name resolution service (for example, the DNS) from being invoked.

13956 If the AI\_NUMERICSERV flag is specified, then a non-null *servname* string supplied shall be a  
13957 numeric port string. Otherwise, an [EAI\_NONAME] error shall be returned. This flag shall  
13958 prevent any type of name resolution service (for example, NIS+) from being invoked.

13959 IP6 If the AI\_V4MAPPED flag is specified along with an *ai\_family* of AF\_INET6, then *getaddrinfo()*  
13960 shall return IPv4-mapped IPv6 addresses on finding no matching IPv6 addresses (*ai\_addrlen*  
13961 shall be 16). The AI\_V4MAPPED flag shall be ignored unless *ai\_family* equals AF\_INET6. If the  
13962 AI\_ALL flag is used with the AI\_V4MAPPED flag, then *getaddrinfo()* shall return all matching  
13963 IPv6 and IPv4 addresses. The AI\_ALL flag without the AI\_V4MAPPED flag is ignored.

13964 The *ai\_socktype* field to which argument *hints* points specifies the socket type for the service, as  
13965 defined in *socket()*. If a specific socket type is not given (for example, a value of zero) and the  
13966 service name could be interpreted as valid with multiple supported socket types, the  
13967 implementation shall attempt to resolve the service name for all supported socket types and, in  
13968 the absence of errors, all possible results shall be returned. A non-zero socket type value shall  
13969 limit the returned information to values with the specified socket type.

13970 If the *ai\_family* field to which *hints* points has the value AF\_UNSPEC, addresses shall be  
13971 returned for use with any address family that can be used with the specified *nodename* and/or  
13972 *servname*. Otherwise, addresses shall be returned for use only with the specified address family.  
13973 If *ai\_family* is not AF\_UNSPEC and *ai\_protocol* is not zero, then addresses are returned for use  
13974 only with the specified address family and protocol; the value of *ai\_protocol* shall be interpreted  
13975 as in a call to the *socket()* function with the corresponding values of *ai\_family* and *ai\_protocol*.

#### 13976 RETURN VALUE

13977 A zero return value for *getaddrinfo()* indicates successful completion; a non-zero return value  
13978 indicates failure. The possible values for the failures are listed in the ERRORS section.

13979 Upon successful return of *getaddrinfo()*, the location to which *res* points shall refer to a linked list  
13980 of **addrinfo** structures, each of which shall specify a socket address and information for use in  
13981 creating a socket with which to use that socket address. The list shall include at least one  
13982 **addrinfo** structure. The *ai\_next* field of each structure contains a pointer to the next structure on  
13983 the list, or a null pointer if it is the last structure on the list. Each structure on the list shall

13984 include values for use with a call to the *socket()* function, and a socket address for use with the  
 13985 *connect()* function or, if the AI\_PASSIVE flag was specified, for use with the *bind()* function. The  
 13986 fields *ai\_family*, *ai\_socktype*, and *ai\_protocol* shall be usable as the arguments to the *socket()*  
 13987 function to create a socket suitable for use with the returned address. The fields *ai\_addr* and  
 13988 *ai\_addrlen* are usable as the arguments to the *connect()* or *bind()* functions with such a socket,  
 13989 according to the AI\_PASSIVE flag.

13990 If *nodename* is not null, and if requested by the AI\_CANONNAME flag, the *ai\_canonname* field of  
 13991 the first returned **addrinfo** structure shall point to a null-terminated string containing the  
 13992 canonical name corresponding to the input *nodename*; if the canonical name is not available, then  
 13993 *ai\_canonname* shall refer to the *nodename* argument or a string with the same contents. The  
 13994 contents of the *ai\_flags* field of the returned structures are undefined.

13995 All fields in socket address structures returned by *getaddrinfo()* that are not filled in through an  
 13996 explicit argument (for example, *sin6\_flowinfo*) shall be set to zero.

13997 **Note:** This makes it easier to compare socket address structures.

#### 13998 **ERRORS**

13999 The *getaddrinfo()* function shall fail and return the corresponding value if:

14000 [EAI\_AGAIN] The name could not be resolved at this time. Future attempts may succeed.

14001 [EAI\_BADFLAGS]

14002 The *flags* parameter had an invalid value.

14003 [EAI\_FAIL] A non-recoverable error occurred when attempting to resolve the name.

14004 [EAI\_FAMILY] The address family was not recognized.

14005 [EAI\_MEMORY] There was a memory allocation failure when trying to allocate storage for the  
 14006 return value.

14007 [EAI\_NONAME] The name does not resolve for the supplied parameters.

14008 Neither *nodename* nor *servname* were supplied. At least one of these shall be  
 14009 supplied.

14010 [EAI\_SERVICE] The service passed was not recognized for the specified socket type.

14011 [EAI\_SOCKTYPE]

14012 The intended socket type was not recognized.

14013 [EAI\_SYSTEM] A system error occurred; the error code can be found in *errno*.

14014 [EAI\_OVERFLOW] An argument buffer overflowed.

#### 14015 **EXAMPLES**

14016 None.

#### 14017 **APPLICATION USAGE**

14018 If the caller handles only TCP and not UDP, for example, then the *ai\_protocol* member of the *hints*  
 14019 structure should be set to IPPROTO\_TCP when *getaddrinfo()* is called.

14020 If the caller handles only IPv4 and not IPv6, then the *ai\_family* member of the *hints* structure  
 14021 should be set to AF\_INET when *getaddrinfo()* is called.

#### 14022 **RATIONALE**

14023 None.

14024 **FUTURE DIRECTIONS**

14025           None.

14026 **SEE ALSO**

14027           *connect()*, *gai\_strerror()*, *gethostbyname()*, *getnameinfo()*, *getservbyname()*, *socket()*, the Base  
14028           Definitions volume of IEEE Std 1003.1-200x, <**netdb.h**>, <**sys/socket.h**>

14029 **CHANGE HISTORY**

14030           First released in Issue 6. Derived from the XNS, Issue 5.2 specification.

14031           The **restrict** keyword is added to the *getaddrinfo()* prototype for alignment with the  
14032           ISO/IEC 9899:1999 standard.

## 14033 NAME

14034 freopen — open a stream

## 14035 SYNOPSIS

14036 #include &lt;stdio.h&gt;

```
14037 FILE *freopen(const char *restrict filename, const char *restrict mode,
14038 FILE *restrict stream);
```

## 14039 DESCRIPTION

14040 CX The functionality described on this reference page is aligned with the ISO C standard. Any  
 14041 conflict between the requirements described here and the ISO C standard is unintentional. This  
 14042 volume of IEEE Std 1003.1-200x defers to the ISO C standard.

14043 The *freopen()* function shall first attempt to flush the stream and close any file descriptor  
 14044 associated with *stream*. Failure to flush or close the file descriptor successfully shall be ignored.  
 14045 The error and end-of-file indicators for the stream shall be cleared.

14046 The *freopen()* function shall open the file whose pathname is the string pointed to by *filename* and  
 14047 associate the stream pointed to by *stream* with it. The *mode* argument shall be used just as in  
 14048 *fopen()*.

14049 The original stream shall be closed regardless of whether the subsequent open succeeds.

14050 If *filename* is a null pointer, the *freopen()* function shall attempt to change the mode of the stream  
 14051 to that specified by *mode*, as if the name of the file currently associated with the stream had been  
 14052 used. It is implementation-defined which changes of mode are permitted (if any), and under  
 14053 what circumstances.

14054 XSI After a successful call to the *freopen()* function, the orientation of the stream shall be cleared, the  
 14055 encoding rule shall be cleared, and the associated **mbstate\_t** object shall be set to describe an  
 14056 initial conversion state.

14057 CX The largest value that can be represented correctly in an object of type **off\_t** shall be established  
 14058 as the offset maximum in the open file description.

## 14059 RETURN VALUE

14060 Upon successful completion, *freopen()* shall return the value of *stream*. Otherwise, a null pointer  
 14061 CX shall be returned, and *errno* shall be set to indicate the error.

## 14062 ERRORS

14063 The *freopen()* function shall fail if:

14064 CX [EACCES] Search permission is denied on a component of the path prefix, or the file  
 14065 exists and the permissions specified by *mode* are denied, or the file does not  
 14066 exist and write permission is denied for the parent directory of the file to be  
 14067 created.

14068 CX [EINTR] A signal was caught during *freopen()*.

14069 CX [EISDIR] The named file is a directory and *mode* requires write access.

14070 CX [ELOOP] A loop exists in symbolic links encountered during resolution of the *path*  
 14071 argument.

14072 CX [EMFILE] {OPEN\_MAX} file descriptors are currently open in the calling process.

14073 CX [ENAMETOOLONG]

14074 The length of the *filename* argument exceeds {PATH\_MAX} or a pathname  
 14075 component is longer than {NAME\_MAX}.

14076 CX	[ENFILE]	The maximum allowable number of files is currently open in the system.
14077 CX 14078	[ENOENT]	A component of <i>filename</i> does not name an existing file or <i>filename</i> is an empty string.
14079 CX 14080	[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.
14081 CX	[ENOTDIR]	A component of the path prefix is not a directory.
14082 CX 14083	[ENXIO]	The named file is a character special or block special file, and the device associated with this special file does not exist.
14084 CX 14085	[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> .
14086 CX 14087	[EROFS]	The named file resides on a read-only file system and <i>mode</i> requires write access.
14088		The <i>freopen()</i> function may fail if:
14089 CX	[EINVAL]	The value of the <i>mode</i> argument is not valid.
14090 CX 14091	[ELOOP]	More than {SYMLOOP_MAX} symbolic links were encountered during resolution of the <i>path</i> argument.
14092 CX 14093 14094	[ENAMETOOLONG]	Pathname resolution of a symbolic link produced an intermediate result whose length exceeds {PATH_MAX}.
14095 CX	[ENOMEM]	Insufficient storage space is available.
14096 CX 14097	[ENXIO]	A request was made of a nonexistent device, or the request was outside the capabilities of the device.
14098 CX 14099	[ETXTBSY]	The file is a pure procedure (shared text) file that is being executed and <i>mode</i> requires write access.

#### 14100 EXAMPLES

##### 14101 Directing Standard Output to a File

14102 The following example logs all standard output to the `/tmp/logfile` file.

```

14103 #include <stdio.h>
14104 ...
14105 FILE *fp;
14106 ...
14107 fp = freopen ("/tmp/logfile", "a+", stdout);
14108 ...

```

#### 14109 APPLICATION USAGE

14110 The *freopen()* function is typically used to attach the preopened *streams* associated with *stdin*,  
 14111 *stdout*, and *stderr* to other files.

#### 14112 RATIONALE

14113 None.

14114 **FUTURE DIRECTIONS**

14115 None.

14116 **SEE ALSO**14117 *fclose()*, *fopen()*, *fdopen()*, *mbsinit()*, the Base Definitions volume of IEEE Std 1003.1-200x,  
14118 **<stdio.h>**14119 **CHANGE HISTORY**

14120 First released in Issue 1. Derived from Issue 1 of the SVID.

14121 **Issue 5**14122 The DESCRIPTION is updated to indicate that the orientation of the stream is cleared and the  
14123 conversion state of the stream is set to an initial conversion state by a successful call to the  
14124 *freopen()* function.

14125 Large File Summit extensions are added.

14126 **Issue 6**

14127 Extensions beyond the ISO C standard are now marked.

14128 The following new requirements on POSIX implementations derive from alignment with the |  
14129 Single UNIX Specification:14130 • In the DESCRIPTION, text is added to indicate setting of the offset maximum in the open file  
14131 description. This change is to support large files.14132 • In the ERRORS section, the [E\_OVERFLOW] condition is added. This change is to support  
14133 large files.

14134 • The [ELOOP] mandatory error condition is added.

14135 • A second [ENAMETOOLONG] is added as an optional error condition.

14136 • The [EINVAL], [ENOMEM], [ENXIO], and [ETXTBSY] optional error conditions are added.

14137 The following changes are made for alignment with the ISO/IEC 9899:1999 standard:

14138 • The *freopen()* prototype is updated.

14139 • The DESCRIPTION is updated.

14140 The wording of the mandatory [ELOOP] error condition is updated, and a second optional  
14141 [ELOOP] error condition is added.

14142 The DESCRIPTION is updated regarding failure to close, changing the “file” to “file descriptor”.

14143 **NAME**

14144 frexp, frexpf, frexpl — extract mantissa and exponent from a double precision number

14145 **SYNOPSIS**

14146 #include <math.h>

14147 double frexp(double *num*, int \**exp*);

14148 float frexpf(float *num*, int \**exp*);

14149 long double frexpl(long double *num*, int \**exp*);

14150 **DESCRIPTION**

14151 **CX** The functionality described on this reference page is aligned with the ISO C standard. Any  
14152 conflict between the requirements described here and the ISO C standard is unintentional. This  
14153 volume of IEEE Std 1003.1-200x defers to the ISO C standard.

14154 These functions shall break a floating-point number *num* into a normalized fraction and an  
14155 integral power of 2. The integer exponent shall be stored in the **int** object pointed to by *exp*.

14156 **RETURN VALUE**

14157 For finite arguments, these functions shall return the value *x*, such that *x* has a magnitude in the  
14158 interval  $[\frac{1}{2}, 1)$  or 0, and *num* equals *x* times 2 raised to the power \**exp*.

14159 **MX** If *num* is NaN, a NaN shall be returned, and the value of \**exp* is unspecified.

14160 If *num* is  $\pm 0$ ,  $\pm 0$  shall be returned, and the value of \**exp* shall be 0.

14161 If *num* is  $\pm \text{Inf}$ , *num* shall be returned, and the value of \**exp* is unspecified.

14162 **ERRORS**

14163 No errors are defined.

14164 **EXAMPLES**

14165 None.

14166 **APPLICATION USAGE**

14167 None.

14168 **RATIONALE**

14169 None.

14170 **FUTURE DIRECTIONS**

14171 None.

14172 **SEE ALSO**

14173 *isnan()*, *ldexp()*, *modf()*, the Base Definitions volume of IEEE Std 1003.1-200x, <math.h>

14174 **CHANGE HISTORY**

14175 First released in Issue 1. Derived from Issue 1 of the SVID.

14176 **Issue 5**

14177 The DESCRIPTION is updated to indicate how an application should check for an error. This  
14178 text was previously published in the APPLICATION USAGE section.

14179 **Issue 6**

14180 The *frexpf()* and *frexpl()* functions are added for alignment with the ISO/IEC 9899:1999  
14181 standard.



14182 The DESCRIPTION, RETURN VALUE, ERRORS, and APPLICATION USAGE sections are |  
14183 revised to align with the ISO/IEC 9899:1999 standard.  
14184 IEC 60559:1989 standard floating-point extensions over the ISO/IEC 9899:1999 standard are  
14185 marked.

## 14186 NAME

14187 fscanf, scanf, sscanf — convert formatted input

## 14188 SYNOPSIS

14189 #include &lt;stdio.h&gt;

14190 int fscanf(FILE \*restrict *stream*, const char \*restrict *format*, ... );14191 int scanf(const char \*restrict *format*, ... );14192 int sscanf(const char \*restrict *s*, const char \*restrict *format*, ... );

## 14193 DESCRIPTION

14194 CX The functionality described on this reference page is aligned with the ISO C standard. Any  
 14195 conflict between the requirements described here and the ISO C standard is unintentional. This  
 14196 volume of IEEE Std 1003.1-200x defers to the ISO C standard.

14197 The *fscanf()* function shall read from the named input *stream*. The *scanf()* function shall read  
 14198 from the standard input stream *stdin*. The *sscanf()* function shall read from the string *s*. Each  
 14199 function reads bytes, interprets them according to a format, and stores the results in its  
 14200 arguments. Each expects, as arguments, a control string *format* described below, and a set of  
 14201 *pointer* arguments indicating where the converted input should be stored. The result is  
 14202 undefined if there are insufficient arguments for the format. If the format is exhausted while  
 14203 arguments remain, the excess arguments shall be evaluated but otherwise ignored.

14204 XSI Conversions can be applied to the *n*th argument after the *format* in the argument list, rather than  
 14205 to the next unused argument. In this case, the conversion specifier character % (see below) is  
 14206 replaced by the sequence "%n\$", where *n* is a decimal integer in the range [1,{NL\_ARGMAX}].  
 14207 This feature provides for the definition of format strings that select arguments in an order  
 14208 appropriate to specific languages. In format strings containing the "%n\$" form of conversion  
 14209 specifications, it is unspecified whether numbered arguments in the argument list can be  
 14210 referenced from the format string more than once.

14211 The *format* can contain either form of a conversion specification—that is, % or "%n\$"—but the  
 14212 two forms cannot be mixed within a single *format* string. The only exception to this is that %% or  
 14213 %\* can be mixed with the "%n\$" form. When numbered argument specifications are used,  
 14214 specifying the *N*th argument requires that all the leading arguments, from the first to the  
 14215 (*N*–1)th, are pointers.

14216 CX The *fscanf()* function in all its forms shall allow detection of a language-dependent radix  
 14217 character in the input string. The radix character is defined in the program's locale (category  
 14218 *LC\_NUMERIC*). In the POSIX locale, or in a locale where the radix character is not defined, the  
 14219 radix character shall default to a period ( '.' ).

14220 The format is a character string, beginning and ending in its initial shift state, if any, composed  
 14221 of zero or more directives. Each directive is composed of one of the following: one or more  
 14222 white-space characters (<space>*s*, <tab>*s*, <newline>*s*, <vertical-tab>*s*, or <form-feed>*s*); an  
 14223 ordinary character (neither '%' nor a white-space character); or a conversion specification. Each  
 14224 XSI conversion specification is introduced by the character '%' or the character sequence "%n\$"  
 14225 after which the following appear in sequence:

- 14226 • An optional assignment-suppressing character '\*'.
- 14227 • An optional non-zero decimal integer that specifies the maximum field width.
- 14228 • An option length modifier that specifies the size of the receiving object.
- 14229 • A *conversion specifier* character that specifies the type of conversion to be applied. The valid  
 14230 conversion specifiers are described below.

14231 The *fscanf()* functions shall execute each directive of the format in turn. If a directive fails, as  
 14232 detailed below, the function shall return. Failures are described as input failures (due to the  
 14233 unavailability of input bytes) or matching failures (due to inappropriate input).

14234 A directive composed of one or more white-space characters shall be executed by reading input  
 14235 until no more valid input can be read, or up to the first byte which is not a white-space character,  
 14236 which remains unread.

14237 A directive that is an ordinary character shall be executed as follows: the next byte shall be read  
 14238 from the input and compared with the byte that comprises the directive; if the comparison  
 14239 shows that they are not equivalent, the directive shall fail, and the differing and subsequent  
 14240 bytes shall remain unread. Similarly, if end-of-file, an encoding error, or a read error prevents a  
 14241 character from being read, the directive shall fail.

14242 A directive that is a conversion specification defines a set of matching input sequences, as  
 14243 described below for each conversion character. A conversion specification shall be executed in  
 14244 the following steps.

14245 Input white-space characters (as specified by *isspace()*) shall be skipped, unless the conversion  
 14246 specification includes a `[`, `c`, `C`, or `n` conversion specifier.

14247 An item shall be read from the input, unless the conversion specification includes an `n`  
 14248 conversion specifier. An input item shall be defined as the longest sequence of input bytes (up to  
 14249 any specified maximum field width, which may be measured in characters or bytes dependent  
 14250 on the conversion specifier) which is an initial subsequence of a matching sequence. The first  
 14251 byte, if any, after the input item shall remain unread. If the length of the input item is 0, the  
 14252 execution of the conversion specification shall fail; this condition is a matching failure, unless  
 14253 end-of-file, an encoding error, or a read error prevented input from the stream, in which case it is  
 14254 an input failure.

14255 Except in the case of a `%` conversion specifier, the input item (or, in the case of a `%n` conversion  
 14256 specification, the count of input bytes) shall be converted to a type appropriate to the conversion  
 14257 character. If the input item is not a matching sequence, the execution of the conversion  
 14258 specification fails; this condition is a matching failure. Unless assignment suppression was  
 14259 indicated by a `'*'`, the result of the conversion shall be placed in the object pointed to by the  
 14260 first argument following the *format* argument that has not already received a conversion result if  
 14261 XSI the conversion specification is introduced by `%`, or in the *n*th argument if introduced by the  
 14262 character sequence `"%n$"`. If this object does not have an appropriate type, or if the result of the  
 14263 conversion cannot be represented in the space provided, the behavior is undefined.

14264 The length modifiers and their meanings are:

14265 `hh` Specifies that a following `d`, `i`, `o`, `u`, `x`, `X`, or `n` conversion specifier applies to an  
 14266 argument with type pointer to **signed char** or **unsigned char**.

14267 `h` Specifies that a following `d`, `i`, `o`, `u`, `x`, `X`, or `n` conversion specifier applies to an  
 14268 argument with type pointer to **short** or **unsigned short**.

14269 `l` (`ell`) Specifies that a following `d`, `i`, `o`, `u`, `x`, `X`, or `n` conversion specifier applies to an  
 14270 argument with type pointer to **long** or **unsigned long**; that a following `a`, `A`, `e`, `E`, `f`, `F`, `g`,  
 14271 or `G` conversion specifier applies to an argument with type pointer to **double**; or that a  
 14272 following `c`, `s`, or `[` conversion specifier applies to an argument with type pointer to  
 14273 **wchar\_t**.

14274 `ll` (`ell-ell`)  
 14275 Specifies that a following `d`, `i`, `o`, `u`, `x`, `X`, or `n` conversion specifier applies to an  
 14276 argument with type pointer to **long long** or **unsigned long long**.

14277	j	Specifies that a following d, i, o, u, x, X, or n conversion specifier applies to an argument with type pointer to <b>intmax_t</b> or <b>uintmax_t</b> .
14278		
14279	z	Specifies that a following d, i, o, u, x, X, or n conversion specifier applies to an argument with type pointer to <b>size_t</b> or the corresponding signed integer type.
14280		
14281	t	Specifies that a following d, i, o, u, x, X, or n conversion specifier applies to an argument with type pointer to <b>ptrdiff_t</b> or the corresponding <b>unsigned</b> type.
14282		
14283	L	Specifies that a following a, A, e, E, f, F, g, or G conversion specifier applies to an argument with type pointer to <b>long double</b> .
14284		
14285		If a length modifier appears with any conversion specifier other than as specified above, the behavior is undefined.
14286		
14287		The following conversion specifiers are valid:
14288	d	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 <b>int</b> .
14289		
14290		
14291		
14292	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 <b>int</b> .
14293		
14294		
14295		
14296	o	Matches an optionally signed octal integer, whose format is the same as expected for the subject sequence of <i>strtol()</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> .
14297		
14298		
14299		
14300	u	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 <b>unsigned</b> .
14301		
14302		
14303		
14304	x	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 <b>unsigned</b> .
14305		
14306		
14307		
14308	a, e, f, g	
14309		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 <b>float</b> .
14310		
14311		
14312		
14313		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.
14314		
14315		
14316	s	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 <b>char</b> , <b>signed char</b> , or <b>unsigned char</b> large enough to accept the sequence and a terminating null character code, which shall be added automatically.
14317		
14318		
14319		
14320		If an l (ell) qualifier is present, the input is a sequence of characters that begins in the initial shift state. Each character shall be converted to a wide character as if by a call to
14321		

14322		the <i>mbrtowc()</i> function, with the conversion state described by an <b>mbstate_t</b> object	
14323		initialized to zero before the first character is converted. The application shall ensure	
14324		that the corresponding argument is a pointer to an array of <b>wchar_t</b> large enough to	
14325		accept the sequence and the terminating null wide character, which shall be added	
14326		automatically.	
14327	[	Matches a non-empty sequence of bytes from a set of expected bytes (the <i>scanset</i> ). The	
14328		normal skip over white-space characters shall be suppressed in this case. The	
14329		application shall ensure that the corresponding argument is a pointer to the initial byte	
14330		of an array of <b>char</b> , <b>signed char</b> , or <b>unsigned char</b> large enough to accept the sequence	
14331		and a terminating null byte, which shall be added automatically.	
14332		If an <b>l</b> (ell) qualifier is present, the input is a sequence of characters that begins in the	
14333		initial shift state. Each character in the sequence shall be converted to a wide character	
14334		as if by a call to the <i>mbrtowc()</i> function, with the conversion state described by an	
14335		<b>mbstate_t</b> object initialized to zero before the first character is converted. The	
14336		application shall ensure that the corresponding argument is a pointer to an array of	
14337		<b>wchar_t</b> large enough to accept the sequence and the terminating null wide character,	
14338		which shall be added automatically.	
14339		The conversion specification includes all subsequent bytes in the <i>format</i> string up to	
14340		and including the matching right square bracket (']'). The bytes between the square	
14341		brackets (the <i>scanlist</i> ) comprise the scanset, unless the byte after the left square bracket	
14342		is a circumflex (^), in which case the scanset contains all bytes that do not appear in	
14343		the scanlist between the circumflex and the right square bracket. If the conversion	
14344		specification begins with "[ ]" or "[ ^ ]", the right square bracket is included in the	
14345		scanlist and the next right square bracket is the matching right square bracket that ends	
14346		the conversion specification; otherwise, the first right square bracket is the one that	
14347		ends the conversion specification. If a '-' is in the scanlist and is not the first character,	
14348		nor the second where the first character is a '^', nor the last character, the behavior is	
14349		implementation-defined.	
14350	c	Matches a sequence of bytes of the number specified by the field width (1 if no field	
14351		width is present in the conversion specification). The application shall ensure that the	
14352		corresponding argument is a pointer to the initial byte of an array of <b>char</b> , <b>signed char</b> ,	
14353		or <b>unsigned char</b> large enough to accept the sequence. No null byte is added. The	
14354		normal skip over white-space characters shall be suppressed in this case.	
14355		If an <b>l</b> (ell) qualifier is present, the input shall be a sequence of characters that begins in	
14356		the initial shift state. Each character in the sequence is converted to a wide character as	
14357		if by a call to the <i>mbrtowc()</i> function, with the conversion state described by an	
14358		<b>mbstate_t</b> object initialized to zero before the first character is converted. The	
14359		application shall ensure that the corresponding argument is a pointer to an array of	
14360		<b>wchar_t</b> large enough to accept the resulting sequence of wide characters. No null wide	
14361		character is added.	
14362	p	Matches an implementation-defined set of sequences, which shall be the same as the set	
14363		of sequences that is produced by the %p conversion specification of the corresponding	
14364		<i>fprintf()</i> functions. The application shall ensure that the corresponding argument is a	
14365		pointer to a pointer to <b>void</b> . The interpretation of the input item is implementation-	
14366		defined. If the input item is a value converted earlier during the same program	
14367		execution, the pointer that results shall compare equal to that value; otherwise, the	
14368		behavior of the %p conversion specification is undefined.	
14369	n	No input is consumed. The application shall ensure that the corresponding argument is	
14370		a pointer to the integer into which shall be written the number of bytes read from the	

14371		input so far by this call to the <i>fscanf()</i> functions. Execution of a <code>%n</code> conversion	
14372		specification shall not increment the assignment count returned at the completion of	
14373		execution of the function. No argument shall be converted, but one shall be consumed.	
14374		If the conversion specification includes an assignment-suppressing character or a field	
14375		width, the behavior is undefined.	
14376	XSI	<b>C</b> Equivalent to <code>lc</code> .	
14377	XSI	<b>S</b> Equivalent to <code>ls</code> .	
14378		<code>%</code> Matches a single <code>'%'</code> character; no conversion or assignment occurs. The complete	
14379		conversion specification shall be <code>%%</code> .	
14380		If a conversion specification is invalid, the behavior is undefined.	
14381		The conversion specifiers <code>A</code> , <code>E</code> , <code>F</code> , <code>G</code> , and <code>X</code> are also valid and shall be equivalent to <code>a</code> , <code>e</code> , <code>f</code> , <code>g</code> , and	
14382		<code>x</code> , respectively.	
14383		If end-of-file is encountered during input, conversion shall be terminated. If end-of-file occurs	
14384		before any bytes matching the current conversion specification (except for <code>%n</code> ) have been read	
14385		(other than leading white-space characters, where permitted), execution of the current	
14386		conversion specification shall terminate with an input failure. Otherwise, unless execution of the	
14387		current conversion specification is terminated with a matching failure, execution of the	
14388		following conversion specification (if any) shall be terminated with an input failure.	
14389		Reaching the end of the string in <i>sscanf()</i> shall be equivalent to encountering end-of-file for	
14390		<i>fscanf()</i> .	
14391		If conversion terminates on a conflicting input, the offending input is left unread in the input.	
14392		Any trailing white space (including <code>&lt;newline&gt;</code> s) shall be left unread unless matched by a	
14393		conversion specification. The success of literal matches and suppressed assignments is only	
14394		directly determinable via the <code>%n</code> conversion specification.	
14395	CX	The <i>fscanf()</i> and <i>scanf()</i> functions may mark the <i>st_atime</i> field of the file associated with <i>stream</i>	
14396		for update. The <i>st_atime</i> field shall be marked for update by the first successful execution of	
14397		<i>fgetc()</i> , <i>fgets()</i> , <i>fread()</i> , <i>getc()</i> , <i>getchar()</i> , <i>gets()</i> , <i>fscanf()</i> , or <i>scanf()</i> using <i>stream</i> that returns data	
14398		not supplied by a prior call to <i>ungetc()</i> .	
14399		<b>RETURN VALUE</b>	
14400		Upon successful completion, these functions shall return the number of successfully matched	
14401		and assigned input items; this number can be zero in the event of an early matching failure. If	
14402		the input ends before the first matching failure or conversion, EOF shall be returned. If a read	
14403	CX	error occurs, the error indicator for the stream is set, EOF shall be returned, and <i>errno</i> shall be set	
14404		to indicate the error.	
14405		<b>ERRORS</b>	
14406		For the conditions under which the <i>fscanf()</i> functions fail and may fail, refer to <i>fgetc()</i> or	
14407		<i>fgetwc()</i> .	
14408		In addition, <i>fscanf()</i> may fail if:	
14409	XSI	<b>[EILSEQ]</b> Input byte sequence does not form a valid character.	
14410	XSI	<b>[EINVAL]</b> There are insufficient arguments.	

14411 **EXAMPLES**

14412 The call:

```
14413 int i, n; float x; char name[50];
14414 n = scanf("%d%f%s", &i, &x, name);
```

14415 with the input line:

14416 25 54.32E-1 Hamster

14417 assigns to *n* the value 3, to *i* the value 25, to *x* the value 5.432, and *name* contains the string  
14418 "Hamster".

14419 The call:

```
14420 int i; float x; char name[50];
14421 (void) scanf("%2d%f*d %[0123456789]", &i, &x, name);
```

14422 with input:

14423 56789 0123 56a72

14424 assigns 56 to *i*, 789.0 to *x*, skips 0123, and places the string "56\0" in *name*. The next call to  
14425 *getchar()* shall return the character 'a'.

14426 **Reading Data into an Array**

14427 The following call uses *fscanf()* to read three floating-point numbers from standard input into  
14428 the *input* array.

```
14429 float input[3]; fscanf (stdin, "%f %f %f", input, input+1, input+2);
```

14430 **APPLICATION USAGE**

14431 If the application calling *fscanf()* has any objects of type **wint\_t** or **wchar\_t**, it must also include  
14432 the **<wchar.h>** header to have these objects defined.

14433 **RATIONALE**

14434 This function is aligned with the ISO/IEC 9899:1999 standard, and in doing so a few "obvious"  
14435 things were not included. Specifically, the set of characters allowed in a scanset is limited to  
14436 single-byte characters. In other similar places, multi-byte characters have been permitted, but  
14437 for alignment with the ISO/IEC 9899:1999 standard, it has not been done here. Applications  
14438 needing this could use the corresponding wide-character functions to achieve the desired  
14439 results.

14440 **FUTURE DIRECTIONS**

14441 None.

14442 **SEE ALSO**

14443 *getc()*, *printf()*, *setlocale()*, *strtod()*, *strtol()*, *strtoul()*, *wcrtomb()*, the Base Definitions volume of  
14444 IEEE Std 1003.1-200x, **<langinfo.h>**, **<stdio.h>**, **<wchar.h>**, the Base Definitions volume of  
14445 IEEE Std 1003.1-200x, Chapter 7, Locale

14446 **CHANGE HISTORY**

14447 First released in Issue 1. Derived from Issue 1 of the SVID.

14448 **Issue 5**

14449 Aligned with ISO/IEC 9899:1990/Amendment 1:1995 (E). Specifically, the **l** (ell) qualifier is  
14450 now defined for the **c**, **s**, and **[** conversion specifiers.

14451 The DESCRIPTION is updated to indicate that if infinity and NaN can be generated by the  
14452 *fprintf()* family of functions, then they are recognized by the *fscanf()* family.

14453 **Issue 6**

14454 The Open Group Corrigendum U021/7 and U028/10 are applied. These correct several  
14455 occurrences of “characters” in the text which have been replaced with the term “bytes”.

14456 The DESCRIPTION is updated to avoid use of the term “must” for application requirements.

14457 The following changes are made for alignment with the ISO/IEC 9899:1999 standard:

14458 • The prototypes for *fscanf()*, *scanf()*, and *sscanf()* are updated.

14459 • The DESCRIPTION is updated. |

14460 • The hh, ll, j, t, and z length modifiers are added. |

14461 • The a, A, and F conversion characters are added. |

14462 The DESCRIPTION is updated to use the terms “conversion specifier” and “conversion  
14463 specification” consistently.



14464 **NAME**

14465       fseek, fseeko — reposition a file-position indicator in a stream

14466 **SYNOPSIS**

14467       #include <stdio.h>

14468       int fseek(FILE \*stream, long offset, int whence);

14469 CX     int fseeko(FILE \*stream, off\_t offset, int whence);

14470

14471 **DESCRIPTION**

14472 CX     The functionality described on this reference page is aligned with the ISO C standard. Any  
14473     conflict between the requirements described here and the ISO C standard is unintentional. This  
14474     volume of IEEE Std 1003.1-200x defers to the ISO C standard.

14475     The *fseek()* function shall set the file-position indicator for the stream pointed to by *stream*. If a  
14476     read or write error occurs, the error indicator for the stream shall be set and *fseek()* fails.

14477     The new position, measured in bytes from the beginning of the file, shall be obtained by adding  
14478     *offset* to the position specified by *whence*. The specified point is the beginning of the file for  
14479     SEEK\_SET, the current value of the file-position indicator for SEEK\_CUR, or end-of-file for  
14480     SEEK\_END.

14481     If the stream is to be used with wide-character input/output functions, the application shall  
14482     ensure that *offset* is either 0 or a value returned by an earlier call to *ftell()* on the same stream and  
14483     *whence* is SEEK\_SET.

14484     A successful call to *fseek()* shall clear the end-of-file indicator for the stream and undo any effects  
14485     of *ungetc()* and *ungetwc()* on the same stream. After an *fseek()* call, the next operation on an  
14486     update stream may be either input or output.

14487 CX     If the most recent operation, other than *ftell()*, on a given stream is *flush()*, the file offset in the  
14488     underlying open file description shall be adjusted to reflect the location specified by *fseek()*.

14489     The *fseek()* function shall allow the file-position indicator to be set beyond the end of existing  
14490     data in the file. If data is later written at this point, subsequent reads of data in the gap shall  
14491     return bytes with the value 0 until data is actually written into the gap.

14492     The behavior of *fseek()* on devices which are incapable of seeking is implementation-defined.  
14493     The value of the file offset associated with such a device is undefined.

14494     If the stream is writable and buffered data had not been written to the underlying file, *fseek()*  
14495     shall cause the unwritten data to be written to the file and shall mark the *st\_ctime* and *st\_mtime*  
14496     fields of the file for update.

14497     In a locale with state-dependent encoding, whether *fseek()* restores the stream's shift state is  
14498     implementation-defined.

14499     The *fseeko()* function shall be equivalent to the *fseek()* function except that the *offset* argument is  
14500     of type **off\_t**.

14501 **RETURN VALUE**

14502 CX     The *fseek()* and *fseeko()* functions shall return 0 if they succeed.

14503 CX     Otherwise, they shall return -1 and set *errno* to indicate the error.

14504 **ERRORS**

14505 CX     The *fseek()* and *fseeko()* functions shall fail if, either the *stream* is unbuffered or the *stream*'s  
14506     buffer needed to be flushed, and the call to *fseek()* or *fseeko()* causes an underlying *lseek()* or  
14507     *write()* to be invoked, and:

14508 CX 14509	[EAGAIN]	The O_NONBLOCK flag is set for the file descriptor and the process would be delayed in the write operation.
14510 CX 14511	[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.
14512 CX	[EFBIG]	An attempt was made to write a file that exceeds the maximum file size.
14513 XSI	[EFBIG]	An attempt was made to write a file that exceeds the process' file size limit.
14514 CX 14515	[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.
14516 CX 14517	[EINTR]	The write operation was terminated due to the receipt of a signal, and no data was transferred.
14518 CX 14519	[EINVAL]	The <i>whence</i> argument is invalid. The resulting file-position indicator would be set to a negative value.
14520 CX 14521 14522 14523 14524	[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.
14525 CX	[ENOSPC]	There was no free space remaining on the device containing the file.
14526 CX 14527	[ENXIO]	A request was made of a nonexistent device, or the request was outside the capabilities of the device.
14528 CX 14529	[EOVERFLOW]	For <i>fseek()</i> , the resulting file offset would be a value which cannot be represented correctly in an object of type <b>long</b> .
14530 CX 14531	[EOVERFLOW]	For <i>fseeko()</i> , the resulting file offset would be a value which cannot be represented correctly in an object of type <b>off_t</b> .
14532 CX 14533	[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.
14534 CX	[ESPIPE]	The file descriptor underlying <i>stream</i> is associated with a pipe or FIFO.

14535 **EXAMPLES**

14536 None.

14537 **APPLICATION USAGE**

14538 None.

14539 **RATIONALE**

14540 None.

14541 **FUTURE DIRECTIONS**

14542 None.

14543 **SEE ALSO**

14544 *fopen()*, *fsetpos()*, *ftell()*, *getrlimit()*, *lseek()*, *rewind()*, *ulimit()*, *ungetc()*, *write()*, the Base  
 14545 Definitions volume of IEEE Std 1003.1-200x, <stdio.h>

14546 **CHANGE HISTORY**

14547 First released in Issue 1. Derived from Issue 1 of the SVID.

14548 **Issue 5**

14549 Normative text previously in the APPLICATION USAGE section is moved to the  
14550 DESCRIPTION.

14551 Large File Summit extensions are added.

14552 **Issue 6**

14553 Extensions beyond the ISO C standard are now marked.

14554 The following new requirements on POSIX implementations derive from alignment with the  
14555 Single UNIX Specification:

- 14556 • The *fseeko()* function is added.
- 14557 • The [EFBIG], [EOVERFLOW], and [ENXIO] mandatory error conditions are added.

14558 The following change is incorporated for alignment with the FIPS requirements:

- 14559 • The [EINTR] error is no longer an indication that the implementation does not report partial  
14560 transfers.

14561 The DESCRIPTION is updated to avoid use of the term “must” for application requirements.

14562 The DESCRIPTION is updated to explicitly state that *fseek()* sets the file-position indicator, and  
14563 then on error the error indicator is set and *fseek()* fails. This is for alignment with the  
14564 ISO/IEC 9899:1999 standard.

14565 **NAME**

14566 fsetpos — set current file position

14567 **SYNOPSIS**

14568 #include &lt;stdio.h&gt;

14569 int fsetpos(FILE \*stream, const fpos\_t \*pos);

14570 **DESCRIPTION**

14571 CX The functionality described on this reference page is aligned with the ISO C standard. Any  
 14572 conflict between the requirements described here and the ISO C standard is unintentional. This  
 14573 volume of IEEE Std 1003.1-200x defers to the ISO C standard.

14574 The *fsetpos()* function shall set the file position and state indicators for the stream pointed to by  
 14575 *stream* according to the value of the object pointed to by *pos*, which the application shall ensure  
 14576 is a value obtained from an earlier call to *fgetpos()* on the same stream. If a read or write error  
 14577 occurs, the error indicator for the stream shall be set and *fsetpos()* fails.

14578 A successful call to the *fsetpos()* function shall clear the end-of-file indicator for the stream and  
 14579 undo any effects of *ungetc()* on the same stream. After an *fsetpos()* call, the next operation on an  
 14580 update stream may be either input or output.

14581 CX The behavior of *fsetpos()* on devices which are incapable of seeking is implementation-defined.  
 14582 The value of the file offset associated with such a device is undefined.

14583 **RETURN VALUE**

14584 The *fsetpos()* function shall return 0 if it succeeds; otherwise, it shall return a non-zero value and  
 14585 set *errno* to indicate the error.

14586 **ERRORS**

14587 CX The *fsetpos()* function shall fail if, either the *stream* is unbuffered or the *stream*'s buffer needed to  
 14588 be flushed, and the call to *fsetpos()* causes an underlying *lseek()* or *write()* to be invoked, and:

14589 CX [EAGAIN] The O\_NONBLOCK flag is set for the file descriptor and the process would be  
 14590 delayed in the write operation.

14591 CX [EBADF] The file descriptor underlying the stream file is not open for writing or the  
 14592 stream's buffer needed to be flushed and the file is not open.

14593 CX [EFBIG] An attempt was made to write a file that exceeds the maximum file size.

14594 XSI [EFBIG] An attempt was made to write a file that exceeds the process' file size limit.

14595 CX [EFBIG] The file is a regular file and an attempt was made to write at or beyond the  
 14596 offset maximum associated with the corresponding stream.

14597 CX [EINTR] The write operation was terminated due to the receipt of a signal, and no data  
 14598 was transferred.

14599 CX [EINVAL] The *whence* argument is invalid. The resulting file-position indicator would be  
 14600 set to a negative value.

14601 CX [EIO] A physical I/O error has occurred, or the process is a member of a  
 14602 background process group attempting to perform a *write()* to its controlling  
 14603 terminal, TOSTOP is set, the process is neither ignoring nor blocking  
 14604 SIGTTOU, and the process group of the process is orphaned. This error may  
 14605 also be returned under implementation-defined conditions.

14606 CX [ENOSPC] There was no free space remaining on the device containing the file.

14607 CX [ENXIO] A request was made of a nonexistent device, or the request was outside the  
14608 capabilities of the device.

14609 CX [EPIPE] The file descriptor underlying *stream* is associated with a pipe or FIFO.

14610 CX [EPIPE] An attempt was made to write to a pipe or FIFO that is not open for reading  
14611 by any process; a SIGPIPE signal shall also be sent to the thread.

#### 14612 EXAMPLES

14613 None.

#### 14614 APPLICATION USAGE

14615 None.

#### 14616 RATIONALE

14617 None.

#### 14618 FUTURE DIRECTIONS

14619 None.

#### 14620 SEE ALSO

14621 *fopen()*, *ftell()*, *lseek()*, *rewind()*, *ungetc()*, *write()*, the Base Definitions volume of  
14622 IEEE Std 1003.1-200x, <stdio.h>

#### 14623 CHANGE HISTORY

14624 First released in Issue 4. Derived from the ISO C standard.

#### 14625 Issue 6

14626 Extensions beyond the ISO C standard are now marked.

14627 An additional [EPIPE] error condition is added for sockets.

14628 The DESCRIPTION is updated to avoid use of the term “must” for application requirements.

14629 The DESCRIPTION is updated to clarify that the error indicator is set for the stream on a read or  
14630 write error. This is for alignment with the ISO/IEC 9899:1999 standard.

## 14631 NAME

14632 fstat — get file status

## 14633 SYNOPSIS

14634 #include &lt;sys/stat.h&gt;

14635 int fstat(int *fildev*, struct stat \**buf*);

## 14636 DESCRIPTION

14637 The *fstat()* function shall obtain information about an open file associated with the file  
 14638 descriptor *fildev*, and shall write it to the area pointed to by *buf*.

14639 SHM If *fildev* references a shared memory object, the implementation shall update in the **stat** structure  
 14640 pointed to by the *buf* argument only the *st\_uid*, *st\_gid*, *st\_size*, and *st\_mode* fields, and only the  
 14641 S\_IRUSR, S\_IWUSR, S\_IRGRP, S\_IWGRP, S\_IROTH, and S\_IWOTH file permission bits need be  
 14642 valid. The implementation may update other fields and flags.

14643 TYM If *fildev* references a typed memory object, the implementation shall update in the **stat** structure  
 14644 pointed to by the *buf* argument only the *st\_uid*, *st\_gid*, *st\_size*, and *st\_mode* fields, and only the  
 14645 S\_IRUSR, S\_IWUSR, S\_IRGRP, S\_IWGRP, S\_IROTH, and S\_IWOTH file permission bits need be  
 14646 valid. The implementation may update other fields and flags.

14647 The *buf* argument is a pointer to a **stat** structure, as defined in <sys/stat.h>, into which  
 14648 information is placed concerning the file.

14649 The structure members *st\_mode*, *st\_ino*, *st\_dev*, *st\_uid*, *st\_gid*, *st\_atime*, *st\_ctime*, and *st\_mtime*  
 14650 shall have meaningful values for all other file types defined in this volume of  
 14651 IEEE Std 1003.1-200x. The value of the member *st\_nlink* shall be set to the number of links to the  
 14652 file.

14653 An implementation that provides additional or alternative file access control mechanisms may,  
 14654 under implementation-defined conditions, cause *fstat()* to fail.

14655 The *fstat()* function shall update any time-related fields as described in the Base Definitions  
 14656 volume of IEEE Std 1003.1-200x, Section 4.7, File Times Update, before writing into the **stat**  
 14657 structure.

## 14658 RETURN VALUE

14659 Upon successful completion, 0 shall be returned. Otherwise, -1 shall be returned and *errno* set to  
 14660 indicate the error.

## 14661 ERRORS

14662 The *fstat()* function shall fail if:

14663 [EBADF] The *fildev* argument is not a valid file descriptor.

14664 [EIO] An I/O error occurred while reading from the file system.

14665 [EOVERFLOW] The file size in bytes or the number of blocks allocated to the file or the file  
 14666 serial number cannot be represented correctly in the structure pointed to by  
 14667 *buf*.

14668 The *fstat()* function may fail if:

14669 [EOVERFLOW] One of the values is too large to store into the structure pointed to by the *buf*  
 14670 argument.

14671 **EXAMPLES**14672 **Obtaining File Status Information**

14673 The following example shows how to obtain file status information for a file named  
 14674 `/home/cnd/mod1`. The structure variable `buffer` is defined for the `stat` structure. The  
 14675 `/home/cnd/mod1` file is opened with read/write privileges and is passed to the open file  
 14676 descriptor `fildev`.

```
14677 #include <sys/types.h>
14678 #include <sys/stat.h>
14679 #include <fcntl.h>

14680 struct stat buffer;
14681 int      status;
14682 ...
14683 fildev = open("/home/cnd/mod1", O_RDWR);
14684 status = fstat(fildev, &buffer);
```

14685 **APPLICATION USAGE**

14686 None.

14687 **RATIONALE**

14688 None.

14689 **FUTURE DIRECTIONS**

14690 None.

14691 **SEE ALSO**

14692 `lstat()`, `stat()`, the Base Definitions volume of IEEE Std 1003.1-200x, `<sys/stat.h>`, `<sys/types.h>`

14693 **CHANGE HISTORY**

14694 First released in Issue 1. Derived from Issue 1 of the SVID.

14695 **Issue 5**

14696 The DESCRIPTION is updated for alignment with the POSIX Realtime Extension.

14697 Large File Summit extensions are added.

14698 **Issue 6**

14699 In the SYNOPSIS, the optional include of the `<sys/types.h>` header is removed.

14700 The following new requirements on POSIX implementations derive from alignment with the  
 14701 Single UNIX Specification:

- 14702 • The requirement to include `<sys/types.h>` has been removed. Although `<sys/types.h>` was  
 14703 required for conforming implementations of previous POSIX specifications, it was not  
 14704 required for UNIX applications.
- 14705 • The [EIO] mandatory error condition is added.
- 14706 • The [EOVERFLOW] mandatory error condition is added. This change is to support large  
 14707 files.
- 14708 • The [EOVERFLOW] optional error condition is added.

14709 The DESCRIPTION is updated for alignment with IEEE Std 1003.1j-2000 by specifying that  
 14710 shared memory object semantics apply to typed memory objects.

14711 **NAME**

14712 fstatvfs, statvfs — get file system information

14713 **SYNOPSIS**

14714 XSI #include &lt;sys/statvfs.h&gt;

14715 int fstatvfs(int *fildev*, struct statvfs \**buf*);14716 int statvfs(const char \*restrict *path*, struct statvfs \*restrict *buf*);

14717

14718 **DESCRIPTION**14719 The *fstatvfs()* function shall obtain information about the file system containing the file |  
14720 referenced by *fildev*. |14721 The *statvfs()* function shall obtain information about the file system containing the file named by |  
14722 *path*. |14723 For both functions, the *buf* argument is a pointer to a **statvfs** structure that shall be filled. Read, |  
14724 write, or execute permission of the named file is not required. |14725 The following flags can be returned in the *f\_flag* member: |

14726 ST\_RDONLY Read-only file system. |

14727 ST\_NOSUID Setuid/setgid bits ignored by *exec*. |14728 It is unspecified whether all members of the **statvfs** structure have meaningful values on all file |  
14729 systems. |14730 **RETURN VALUE**14731 Upon successful completion, *statvfs()* shall return 0. Otherwise, it shall return  $-1$  and set *errno* to |  
14732 indicate the error. |14733 **ERRORS**14734 The *fstatvfs()* and *statvfs()* functions shall fail if:

14735 [EIO] An I/O error occurred while reading the file system.

14736 [EINTR] A signal was caught during execution of the function.

14737 [EOVERFLOW] One of the values to be returned cannot be represented correctly in the |  
14738 structure pointed to by *buf*. |14739 The *fstatvfs()* function shall fail if:14740 [EBADF] The *fildev* argument is not an open file descriptor.14741 The *statvfs()* function shall fail if:

14742 [EACCES] Search permission is denied on a component of the path prefix.

14743 [ELOOP] A loop exists in symbolic links encountered during resolution of the *path* |  
14744 argument. |

14745 [ENAMETOOLONG]

14746 The length of a pathname exceeds {PATH\_MAX} or a pathname component is |  
14747 longer than {NAME\_MAX}. |14748 [ENOENT] A component of *path* does not name an existing file or *path* is an empty string.14749 [ENOTDIR] A component of the path prefix of *path* is not a directory.14750 The *statvfs()* function may fail if:



14751 [ELOOP] More than {SYMLOOP\_MAX} symbolic links were encountered during  
14752 resolution of the *path* argument.

14753 [ENAMETOOLONG]  
14754 Pathname resolution of a symbolic link produced an intermediate result |  
14755 whose length exceeds {PATH\_MAX}.

#### 14756 EXAMPLES

##### 14757 Obtaining File System Information Using *fstatvfs()*

14758 The following example shows how to obtain file system information for the file system upon  
14759 which the file named */home/cnd/mod1* resides, using the *fstatvfs()* function. The  
14760 */home/cnd/mod1* file is opened with read/write privileges and the open file descriptor is passed  
14761 to the *fstatvfs()* function.

```
14762 #include <statvfs.h>
14763 #include <fcntl.h>
14764 struct statvfs buffer;
14765 int status;
14766 ...
14767 fildes = open("/home/cnd/mod1", O_RDWR);
14768 status = fstatvfs(fildes, &buffer);
```

##### 14769 Obtaining File System Information Using *statvfs()*

14770 The following example shows how to obtain file system information for the file system upon  
14771 which the file named */home/cnd/mod1* resides, using the *statvfs()* function.

```
14772 #include <statvfs.h>
14773 struct statvfs buffer;
14774 int status;
14775 ...
14776 status = statvfs("/home/cnd/mod1", &buffer);
```

#### 14777 APPLICATION USAGE

14778 None.

#### 14779 RATIONALE

14780 None.

#### 14781 FUTURE DIRECTIONS

14782 None.

#### 14783 SEE ALSO

14784 *chmod()*, *chown()*, *creat()*, *dup()*, *exec*, *fcntl()*, *link()*, *mknod()*, *open()*, *pipe()*, *read()*, *time()*,  
14785 *unlink()*, *utime()*, *write()*, the Base Definitions volume of IEEE Std 1003.1-200x, <*sys/statvfs.h*>

#### 14786 CHANGE HISTORY

14787 First released in Issue 4, Version 2.

#### 14788 Issue 5

14789 Moved from X/OPEN UNIX extension to BASE.

14790 Large File Summit extensions are added.

14791 **Issue 6**

14792 The DESCRIPTION is updated to avoid use of the term “must” for application requirements.

14793 The **restrict** keyword is added to the *statvfs()* prototype for alignment with the  
14794 ISO/IEC 9899:1999 standard.

14795 The wording of the mandatory [ELOOP] error condition is updated, and a second optional  
14796 [ELOOP] error condition is added.

14797 **NAME**

14798 fsync — synchronize changes to a file

14799 **SYNOPSIS**

14800 FSC #include &lt;unistd.h&gt;

14801 int fsync(int *fildev*);

14802

14803 **DESCRIPTION**

14804 The *fsync()* function shall request that all data for the open file descriptor named by *fildev* is to be  
 14805 transferred to the storage device associated with the file described by *fildev* in an  
 14806 implementation-defined manner. The *fsync()* function shall not return until the system has  
 14807 completed that action or until an error is detected.

14808 SIO If `_POSIX_SYNCHRONIZED_IO` is defined, the *fsync()* function shall force all currently queued  
 14809 I/O operations associated with the file indicated by file descriptor *fildev* to the synchronized I/O  
 14810 completion state. All I/O operations shall be completed as defined for synchronized I/O file  
 14811 integrity completion.

14812 **RETURN VALUE**

14813 Upon successful completion, *fsync()* shall return 0. Otherwise, `-1` shall be returned and *errno* set  
 14814 to indicate the error. If the *fsync()* function fails, outstanding I/O operations are not guaranteed  
 14815 to have been completed.

14816 **ERRORS**14817 The *fsync()* function shall fail if:14818 [EBADF] The *fildev* argument is not a valid descriptor.14819 [EINTR] The *fsync()* function was interrupted by a signal.14820 [EINVAL] The *fildev* argument does not refer to a file on which this operation is possible.

14821 [EIO] An I/O error occurred while reading from or writing to the file system.

14822 In the event that any of the queued I/O operations fail, *fsync()* shall return the error conditions  
 14823 defined for *read()* and *write()*.

14824 **EXAMPLES**

14825 None.

14826 **APPLICATION USAGE**

14827 The *fsync()* function should be used by programs which require modifications to a file to be  
 14828 completed before continuing; for example, a program which contains a simple transaction  
 14829 facility might use it to ensure that all modifications to a file or files caused by a transaction are  
 14830 recorded.

14831 **RATIONALE**

14832 The *fsync()* function is intended to force a physical write of data from the buffer cache, and to  
 14833 assure that after a system crash or other failure that all data up to the time of the *fsync()* call is  
 14834 recorded on the disk. Since the concepts of “buffer cache”, “system crash”, “physical write”, and  
 14835 “non-volatile storage” are not defined here, the wording has to be more abstract.

14836 If `_POSIX_SYNCHRONIZED_IO` is not defined, the wording relies heavily on the conformance  
 14837 document to tell the user what can be expected from the system. It is explicitly intended that a  
 14838 null implementation is permitted. This could be valid in the case where the system cannot assure  
 14839 non-volatile storage under any circumstances or when the system is highly fault-tolerant and the  
 14840 functionality is not required. In the middle ground between these extremes, *fsync()* might or  
 14841 might not actually cause data to be written where it is safe from a power failure. The

14842 conformance document should identify at least that one configuration exists (and how to obtain  
14843 that configuration) where this can be assured for at least some files that the user can select to use  
14844 for critical data. It is not intended that an exhaustive list is required, but rather sufficient  
14845 information is provided to let the user determine that if he or she has critical data he or she can  
14846 configure her system to allow it to be written to non-volatile storage.

14847 It is reasonable to assert that the key aspects of *fsync()* are unreasonable to test in a test suite.  
14848 That does not make the function any less valuable, just more difficult to test. A formal  
14849 conformance test should probably force a system crash (power shutdown) during the test for  
14850 this condition, but it needs to be done in such a way that automated testing does not require this  
14851 to be done except when a formal record of the results is being made. It would also not be  
14852 unreasonable to omit testing for *fsync()*, allowing it to be treated as a quality-of-implementation  
14853 issue.

#### 14854 FUTURE DIRECTIONS

14855 None.

#### 14856 SEE ALSO

14857 *sync()*, the Base Definitions volume of IEEE Std 1003.1-200x, <**unistd.h**>

#### 14858 CHANGE HISTORY

14859 First released in Issue 3.

#### 14860 Issue 5

14861 Aligned with *fsync()* in the POSIX Realtime Extension. Specifically, the DESCRIPTION and  
14862 RETURN VALUE sections are much expanded, and the ERRORS section is updated to indicate  
14863 that *fsync()* can return the error conditions defined for *read()* and *write()*.

#### 14864 Issue 6

14865 This function is marked as part of the File Synchronization option.

14866 The following new requirements on POSIX implementations derive from alignment with the  
14867 Single UNIX Specification:

- 14868 • The [EINVAL] and [EIO] mandatory error conditions are added.

14869 **NAME**

14870 ftell, ftello — return a file offset in a stream

14871 **SYNOPSIS**

14872 #include &lt;stdio.h&gt;

14873 long ftell(FILE \*stream);

14874 CX off\_t ftello(FILE \*stream);

14875

14876 **DESCRIPTION**

14877 CX The functionality described on this reference page is aligned with the ISO C standard. Any  
 14878 conflict between the requirements described here and the ISO C standard is unintentional. This  
 14879 volume of IEEE Std 1003.1-200x defers to the ISO C standard.

14880 The *ftell()* function shall obtain the current value of the file-position indicator for the stream  
 14881 pointed to by *stream*.

14882 CX The *ftello()* function shall be equivalent to *ftell()*, except that the return value is of type **off\_t**.

14883 **RETURN VALUE**

14884 CX Upon successful completion, *ftell()* and *ftello()* shall return the current value of the file-position  
 14885 indicator for the stream measured in bytes from the beginning of the file.

14886 CX Otherwise, *ftell()* and *ftello()* shall return  $-1$ , cast to **long** and **off\_t** respectively, and set *errno* to  
 14887 indicate the error.

14888 **ERRORS**

14889 CX The *ftell()* and *ftello()* functions shall fail if:

14890 CX [EBADF] The file descriptor underlying *stream* is not an open file descriptor.

14891 CX [EOVERFLOW] For *ftell()*, the current file offset cannot be represented correctly in an object of  
 14892 type **long**.

14893 CX [EOVERFLOW] For *ftello()*, the current file offset cannot be represented correctly in an object  
 14894 of type **off\_t**.

14895 CX [ESPIPE] The file descriptor underlying *stream* is associated with a pipe or FIFO.

14896 The *ftell()* function may fail if:

14897 CX [ESPIPE] The file descriptor underlying *stream* is associated with a socket.

14898 **EXAMPLES**

14899 None.

14900 **APPLICATION USAGE**

14901 None.

14902 **RATIONALE**

14903 None.

14904 **FUTURE DIRECTIONS**

14905 None.

14906 **SEE ALSO**

14907 *fgetpos()*, *fopen()*, *fseek()*, *lseek()*, the Base Definitions volume of IEEE Std 1003.1-200x, <stdio.h>

14908 **CHANGE HISTORY**

14909 First released in Issue 1. Derived from Issue 1 of the SVID.

14910 **Issue 5**

14911 Large File Summit extensions are added.

14912 **Issue 6**

14913 Extensions beyond the ISO C standard are now marked.

14914 The following new requirements on POSIX implementations derive from alignment with the  
14915 Single UNIX Specification:

- 14916 • The *ftello()* function is added.
- 14917 • The [Eoverflow] error conditions are added.

14918 An additional [ESPIPE] error condition is added for sockets.

14919 **NAME**14920 ftime — get date and time (**LEGACY**)14921 **SYNOPSIS**

14922 xSI #include &lt;sys/timeb.h&gt;

14923 int ftime(struct timeb \*tp);

14924

14925 **DESCRIPTION**

14926 The *ftime()* function shall set the *time* and *millitm* members of the **timeb** structure pointed to by  
 14927 *tp* to contain the seconds and milliseconds portions, respectively, of the current time in seconds  
 14928 since the Epoch. The contents of the *timezone* and *dstflag* members of *tp* after a call to *ftime()* are  
 14929 unspecified.

14930 The system clock need not have millisecond granularity. Depending on any granularity  
 14931 (particularly a granularity of one) renders code non-portable.

14932 **RETURN VALUE**14933 Upon successful completion, the *ftime()* function shall return 0; otherwise, -1 shall be returned.14934 **ERRORS**

14935 No errors are defined.

14936 **EXAMPLES**14937 **Getting the Current Time and Date**

14938 The following example shows how to get the current system time values using the *ftime()*  
 14939 function. The **timeb** structure pointed to by *tp* is filled with the current system time values for  
 14940 *time* and *millitm*.

14941 #include &lt;sys/timeb.h&gt;

14942 struct timeb tp;

14943 int status;

14944 ...

14945 status = ftime(&amp;tp);

14946 **APPLICATION USAGE**

14947 For applications portability, the *time()* function should be used to determine the current time  
 14948 instead of *ftime()*. Realtime applications should use *clock\_gettime()* to determine the current  
 14949 time instead of *ftime()*.

14950 **RATIONALE**

14951 None.

14952 **FUTURE DIRECTIONS**

14953 This function may be withdrawn in a future version.

14954 **SEE ALSO**

14955 *clock\_getres()*, *ctime()*, *gettimeofday()*, *time()*, the Base Definitions volume of  
 14956 IEEE Std 1003.1-200x, <sys/timeb.h>

14957 **CHANGE HISTORY**

14958 First released in Issue 4, Version 2.

14959 **Issue 5**

14960 Moved from X/OPEN UNIX extension to BASE.

14961 Normative text previously in the APPLICATION USAGE section is moved to the  
14962 DESCRIPTION.

14963 **Issue 6**

14964 This function is marked LEGACY.

14965 The DESCRIPTION is updated to refer to “seconds since the Epoch” rather than “seconds since  
14966 00:00:00 UTC (Coordinated Universal Time), January 1 1970” for consistency with other *time*  
14967 functions.



14968 **NAME**

14969 ftok — generate an IPC key

14970 **SYNOPSIS**

14971 xSI #include &lt;sys/ipc.h&gt;

14972 key\_t ftok(const char \*path, int id);

14973

14974 **DESCRIPTION**

14975 The *ftok()* function shall return a key based on *path* and *id* that is usable in subsequent calls to  
 14976 *msgget()*, *semget()*, and *shmget()*. The application shall ensure that the *path* argument is the  
 14977 pathname of an existing file that the process is able to *stat()*.

14978 The *ftok()* function shall return the same key value for all paths that name the same file, when  
 14979 called with the same *id* value, and return different key values when called with different *id*  
 14980 values or with paths that name different files existing on the same file system at the same time. It  
 14981 is unspecified whether *ftok()* shall return the same key value when called again after the file  
 14982 named by *path* is removed and recreated with the same name.

14983 Only the low order 8-bits of *id* are significant. The behavior of *ftok()* is unspecified if these bits  
 14984 are 0.

14985 **RETURN VALUE**

14986 Upon successful completion, *ftok()* shall return a key. Otherwise, *ftok()* shall return (**key\_t**)-1  
 14987 and set *errno* to indicate the error.

14988 **ERRORS**14989 The *ftok()* function shall fail if:

14990 [EACCES] Search permission is denied for a component of the path prefix.

14991 [ELOOP] A loop exists in symbolic links encountered during resolution of the *path*  
 14992 argument.

14993 [ENAMETOOLONG]

14994 The length of the *path* argument exceeds {PATH\_MAX} or a pathname  
 14995 component is longer than {NAME\_MAX}.

14996 [ENOENT] A component of *path* does not name an existing file or *path* is an empty string.

14997 [ENOTDIR] A component of the path prefix is not a directory.

14998 The *ftok()* function may fail if:

14999 [ELOOP] More than {SYMLOOP\_MAX} symbolic links were encountered during  
 15000 resolution of the *path* argument.

15001 [ENAMETOOLONG]

15002 Pathname resolution of a symbolic link produced an intermediate result  
 15003 whose length exceeds {PATH\_MAX}.

15004 **EXAMPLES**15005 **Getting an IPC Key**

15006 The following example gets a unique key that can be used by the IPC functions *semget()*,  
 15007 *msgget()*, and *shmget()*. The key returned by *ftok()* for this example is based on the ID value *S* |  
 15008 and the pathname */tmp*. |

```
15009 #include <sys/ipc.h>
15010 ...
15011 key_t key;
15012 char *path = "/tmp";
15013 int id = 'S';
15014 key = ftok(path, id);
```

15015 **Saving an IPC Key**

15016 The following example gets a unique key based on the pathname */tmp* and the ID value *a*. It |  
 15017 also assigns the value of the resulting key to the *semkey* variable so that it will be available to a  
 15018 later call to *semget()*, *msgget()*, or *shmget()*.

```
15019 #include <sys/ipc.h>
15020 ...
15021 key_t semkey;
15022 if ((semkey = ftok("/tmp", 'a')) == (key_t) -1) {
15023     perror("IPC error: ftok"); exit(1);
15024 }
```

15025 **APPLICATION USAGE**

15026 For maximum portability, *id* should be a single-byte character.

15027 **RATIONALE**

15028 None.

15029 **FUTURE DIRECTIONS**

15030 None.

15031 **SEE ALSO**

15032 *msgget()*, *semget()*, *shmget()*, the Base Definitions volume of IEEE Std 1003.1-200x, *<sys/ipc.h>*

15033 **CHANGE HISTORY**

15034 First released in Issue 4, Version 2.

15035 **Issue 5**

15036 Moved from X/OPEN UNIX extension to BASE.

15037 **Issue 6**

15038 The DESCRIPTION is updated to avoid use of the term “must” for application requirements.

15039 The wording of the mandatory [ELOOP] error condition is updated, and a second optional  
 15040 [ELOOP] error condition is added.

15041 **NAME**

15042        ftruncate — truncate a file to a specified length

15043 **SYNOPSIS**

15044        #include &lt;unistd.h&gt;

15045        int ftruncate(int *fd*, off\_t *length*);15046 **DESCRIPTION**15047        If *fd* is not a valid file descriptor open for writing, the *ftruncate()* function shall fail.

15048        If *fd* 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()*.

15054        Upon successful completion, if *fd* 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.

15057 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 thread.

15059        If *fd* refers to a directory, *ftruncate()* shall fail.

15060        If *fd* refers to any other file type, except a shared memory object, the result is unspecified.

15061 SHM        If *fd* refers to a shared memory object, *ftruncate()* shall set the size of the shared memory object to *length*.

15063 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.

15066 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.

15068 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.

15071 **RETURN VALUE**

15072        Upon successful completion, *ftruncate()* shall return 0; otherwise, -1 shall be returned and *errno* set to indicate the error.

15074 **ERRORS**

15075        The *ftruncate()* function shall fail if:

15076        [EINTR]        A signal was caught during execution.

15077        [EINVAL]       The *length* argument was less than 0.

15078        [EFBIG] or [EINVAL]

15079        The *length* argument was greater than the maximum file size.

15080 XSI        [EFBIG]       The file is a regular file and *length* is greater than the offset maximum established in the open file description associated with *fd*.

15082        [EIO]         An I/O error occurred while reading from or writing to a file system.

- 15083 [EBADF] or [EINVAL]  
15084 The *fildest* argument is not a file descriptor open for writing.
- 15085 [EINVAL]  
15086 The *fildest* argument references a file that was opened without write permission.
- 15087 [EROFS] The named file resides on a read-only file system.

**15088 EXAMPLES**

15089 None.

**15090 APPLICATION USAGE**

15091 None.

**15092 RATIONALE**

15093 The *ftruncate()* function is part of IEEE Std 1003.1-200x as it was deemed to be more useful than  
15094 *truncate()*. The *truncate()* function is provided as an XSI extension.

**15095 FUTURE DIRECTIONS**

15096 None.

**15097 SEE ALSO**

15098 *open()*, *truncate()*, the Base Definitions volume of IEEE Std 1003.1-200x, <**unistd.h**>

**15099 CHANGE HISTORY**

15100 First released in Issue 4, Version 2.

**15101 Issue 5**

15102 Moved from X/OPEN UNIX extension to BASE and aligned with *ftruncate()* in the POSIX  
15103 Realtime Extension. Specifically, the DESCRIPTION is extensively reworded and [EROFS] is  
15104 added to the list of mandatory errors that can be returned by *ftruncate()*.

15105 Large File Summit extensions are added.

**15106 Issue 6**

15107 The *truncate()* function has been split out into a separate reference page.

15108 The following new requirements on POSIX implementations derive from alignment with the  
15109 Single UNIX Specification:

- 15110 • The DESCRIPTION is change to indicate that if the file size is changed, and if the file is a  
15111 regular file, the S\_ISUID and S\_ISGID bits in the file mode may be cleared.

15112 The following changes were made to align with the IEEE P1003.1a draft standard:

- 15113 • The DESCRIPTION text is updated.

15114 XSI-conformant systems are required to increase the size of the file if the file was previously  
15115 smaller than the size requested.

15116 **NAME**

15117       ftrylockfile — stdio locking functions

15118 **SYNOPSIS**

15119 TSF     #include &lt;stdio.h&gt;

15120       int ftrylockfile(FILE \*file);

15121

15122 **DESCRIPTION**15123       Refer to *flockfile()*.

15124 **NAME**

15125 ftw — traverse (walk) a file tree

15126 **SYNOPSIS**

```
15127 xSI #include <ftw.h>
15128 int ftw(const char *path, int (*fn)(const char *,
15129     const struct stat *ptr, int flag), int ndirs);
15130
```

15131 **DESCRIPTION**

15132 The *ftw()* function shall recursively descend the directory hierarchy rooted in *path*. For each  
 15133 object in the hierarchy, *ftw()* shall call the function pointed to by *fn*, passing it a pointer to a  
 15134 null-terminated character string containing the name of the object, a pointer to a **stat** structure  
 15135 containing information about the object, and an integer. Possible values of the integer, defined  
 15136 in the `<ftw.h>` header, are:

15137 **FTW\_D** For a directory.

15138 **FTW\_DNR** For a directory that cannot be read.

15139 **FTW\_F** For a file.

15140 **FTW\_SL** For a symbolic link (but see also **FTW\_NS** below).

15141 **FTW\_NS** For an object other than a symbolic link on which *stat()* could not successfully be  
 15142 executed. If the object is a symbolic link and *stat()* failed, it is unspecified whether  
 15143 *ftw()* passes **FTW\_SL** or **FTW\_NS** to the user-supplied function.

15144 If the integer is **FTW\_DNR**, descendants of that directory shall not be processed. If the integer is  
 15145 **FTW\_NS**, the **stat** structure contains undefined values. An example of an object that would  
 15146 cause **FTW\_NS** to be passed to the function pointed to by *fn* would be a file in a directory with  
 15147 read but without execute (search) permission.

15148 The *ftw()* function shall visit a directory before visiting any of its descendants.

15149 The *ftw()* function shall use at most one file descriptor for each level in the tree.

15150 The argument *ndirs* should be in the range of 1 to `{OPEN_MAX}`.

15151 The tree traversal shall continue until either the tree is exhausted, an invocation of *fn* returns a  
 15152 non-zero value, or some error, other than `[EACCES]`, is detected within *ftw()*.

15153 The *ndirs* argument shall specify the maximum number of directory streams or file descriptors  
 15154 or both available for use by *ftw()* while traversing the tree. When *ftw()* returns it shall close any  
 15155 directory streams and file descriptors it uses not counting any opened by the application-  
 15156 supplied *fn* function.

15157 The results are unspecified if the application-supplied *fn* function does not preserve the current  
 15158 working directory.

15159 The *ftw()* function need not be reentrant. A function that is not required to be reentrant is not  
 15160 required to be thread-safe.

15161 **RETURN VALUE**

15162 If the tree is exhausted, *ftw()* shall return 0. If the function pointed to by *fn* returns a non-zero  
 15163 value, *ftw()* shall stop its tree traversal and return whatever value was returned by the function  
 15164 pointed to by *fn*. If *ftw()* detects an error, it shall return `-1` and set *errno* to indicate the error.

15165 If *ftw()* encounters an error other than `[EACCES]` (see **FTW\_DNR** and **FTW\_NS** above), it shall  
 15166 return `-1` and set *errno* to indicate the error. The external variable *errno* may contain any error

15167 value that is possible when a directory is opened or when one of the *stat* functions is executed on  
15168 a directory or file.

### 15169 ERRORS

15170 The *ftw()* function shall fail if:

15171 [EACCES] Search permission is denied for any component of *path* or read permission is  
15172 denied for *path*.

15173 [ELOOP] A loop exists in symbolic links encountered during resolution of the *path*  
15174 argument.

15175 [ENAMETOOLONG]  
15176 The length of the *path* argument exceeds {PATH\_MAX} or a pathname |  
15177 component is longer than {NAME\_MAX}. |

15178 [ENOENT] A component of *path* does not name an existing file or *path* is an empty string.

15179 [ENOTDIR] A component of *path* is not a directory. |

15180 [EOVERFLOW] A field in the *stat* structure cannot be represented correctly in the current |  
15181 programming environment for one or more files found in the file hierarchy. |

15182 The *ftw()* function may fail if:

15183 [EINVAL] The value of the *ndirs* argument is invalid.

15184 [ELOOP] More than {SYMLOOP\_MAX} symbolic links were encountered during  
15185 resolution of the *path* argument.

15186 [ENAMETOOLONG]  
15187 Pathname resolution of a symbolic link produced an intermediate result |  
15188 whose length exceeds {PATH\_MAX}. |

15189 In addition, if the function pointed to by *fn* encounters system errors, *errno* may be set  
15190 accordingly.

### 15191 EXAMPLES

#### 15192 Walking a Directory Structure

15193 The following example walks the current directory structure, calling the *fn* function for every  
15194 directory entry, using at most 10 file descriptors:

```
15195 #include <ftw.h>
15196 ...
15197 if (ftw(".", fn, 10) != 0) {
15198     perror("ftw"); exit(2);
15199 }
```

### 15200 APPLICATION USAGE

15201 The *ftw()* function may allocate dynamic storage during its operation. If *ftw()* is forcibly  
15202 terminated, such as by *longjmp()* or *siglongjmp()* being executed by the function pointed to by *fn*  
15203 or an interrupt routine, *ftw()* does not have a chance to free that storage, so it remains  
15204 permanently allocated. A safe way to handle interrupts is to store the fact that an interrupt has  
15205 occurred, and arrange to have the function pointed to by *fn* return a non-zero value at its next  
15206 invocation.

15207 **RATIONALE**

15208           None.

15209 **FUTURE DIRECTIONS**

15210           None.

15211 **SEE ALSO**

15212           *longjmp()*, *lstat()*, *malloc()*, *nftw()*, *opendir()*, *siglongjmp()*, *stat()*, the Base Definitions volume of  
15213           IEEE Std 1003.1-200x, <ftw.h>, <sys/stat.h>

15214 **CHANGE HISTORY**

15215           First released in Issue 1. Derived from Issue 1 of the SVID.

15216 **Issue 5**

15217           UX codings in the DESCRIPTION, RETURN VALUE, and ERRORS sections have been changed  
15218           to EX.

15219 **Issue 6**

15220           The ERRORS section is updated as follows: |

15221           • The wording of the mandatory [ELOOP] error condition is updated. |

15222           • A second optional [ELOOP] error condition is added. |

15223           • The [EOVERFLOW] mandatory error condition is added. |

15224           Text is added to the DESCRIPTION to say that the *ftw()* function need not be reentrant and that |  
15225           the results are unspecified if the application-supplied *fn* function does not preserve the current |  
15226           working directory. |



15227 **NAME**

15228 funlockfile — stdio locking functions

15229 **SYNOPSIS**

15230 TSF #include &lt;stdio.h&gt;

15231 void funlockfile(FILE \*file);

15232

15233 **DESCRIPTION**15234 Refer to *flockfile()*.

15235 **NAME**

15236 fwide — set stream orientation

15237 **SYNOPSIS**

15238 #include &lt;stdio.h&gt;

15239 #include &lt;wchar.h&gt;

15240 int fwide(FILE \**stream*, int *mode*);15241 **DESCRIPTION**

15242 CX The functionality described on this reference page is aligned with the ISO C standard. Any  
15243 conflict between the requirements described here and the ISO C standard is unintentional. This  
15244 volume of IEEE Std 1003.1-200x defers to the ISO C standard.

15245 The *fwide()* function shall determine the orientation of the stream pointed to by *stream*. If *mode* is  
15246 greater than zero, the function first attempts to make the stream wide-oriented. If *mode* is less  
15247 than zero, the function first attempts to make the stream byte-oriented. Otherwise, *mode* is zero  
15248 and the function does not alter the orientation of the stream.

15249 If the orientation of the stream has already been determined, *fwide()* shall not change it.

15250 CX Since no return value is reserved to indicate an error, an application wishing to check for error  
15251 situations should set *errno* to 0, then call *fwide()*, then check *errno*, and if it is non-zero, assume  
15252 an error has occurred.

15253 **RETURN VALUE**

15254 The *fwide()* function shall return a value greater than zero if, after the call, the stream has wide-  
15255 orientation, a value less than zero if the stream has byte-orientation, or zero if the stream has no  
15256 orientation.

15257 **ERRORS**

15258 The *fwide()* function may fail if:

15259 CX [EBADF] The *stream* argument is not a valid stream.

15260 **EXAMPLES**

15261 None.

15262 **APPLICATION USAGE**

15263 A call to *fwide()* with *mode* set to zero can be used to determine the current orientation of a  
15264 stream.

15265 **RATIONALE**

15266 None.

15267 **FUTURE DIRECTIONS**

15268 None.

15269 **SEE ALSO**

15270 The Base Definitions volume of IEEE Std 1003.1-200x, <wchar.h>

15271 **CHANGE HISTORY**

15272 First released in Issue 5. Included for alignment with ISO/IEC 9899:1990/Amendment 1:1995  
15273 (E).

15274 **Issue 6**

15275 Extensions beyond the ISO C standard are now marked.

## 15276 NAME

15277 fwprintf, swprintf, wprintf — print formatted wide-character output

## 15278 SYNOPSIS

15279 #include &lt;stdio.h&gt;

15280 #include &lt;wchar.h&gt;

15281 int fwprintf(FILE \*restrict *stream*, const wchar\_t \*restrict *format*, ...);15282 int swprintf(wchar\_t \*restrict *ws*, size\_t *n*,15283 const wchar\_t \*restrict *format*, ...);15284 int wprintf(const wchar\_t \*restrict *format*, ...);

## 15285 DESCRIPTION

15286 CX The functionality described on this reference page is aligned with the ISO C standard. Any  
 15287 conflict between the requirements described here and the ISO C standard is unintentional. This  
 15288 volume of IEEE Std 1003.1-200x defers to the ISO C standard.

15289 The *fwprintf()* function shall place output on the named output *stream*. The *wprintf()* function |  
 15290 shall place output on the standard output stream *stdout*. The *swprintf()* function shall place |  
 15291 output followed by the null wide character in consecutive wide characters starting at *\*ws*; no |  
 15292 more than *n* wide characters shall be written, including a terminating null wide character, which |  
 15293 is always added (unless *n* is zero).

15294 Each of these functions shall convert, format, and print its arguments under control of the *format* |  
 15295 wide-character string. The *format* is composed of zero or more directives: *ordinary wide-* |  
 15296 *characters*, which are simply copied to the output stream, and *conversion specifications*, each of |  
 15297 which results in the fetching of zero or more arguments. The results are undefined if there are |  
 15298 insufficient arguments for the *format*. If the *format* is exhausted while arguments remain, the |  
 15299 excess arguments are evaluated but are otherwise ignored.

15300 XSI Conversions can be applied to the *n*th argument after the *format* in the argument list, rather than |  
 15301 to the next unused argument. In this case, the conversion specifier wide character % (see below) |  
 15302 is replaced by the sequence "%n\$", where *n* is a decimal integer in the range [1,{NL\_ARGMAX}], |  
 15303 giving the position of the argument in the argument list. This feature provides for the definition |  
 15304 of *format* wide-character strings that select arguments in an order appropriate to specific |  
 15305 languages (see the EXAMPLES section).

15306 The *format* can contain either numbered argument specifications (that is, "%n\$" and "\*m\$"), or |  
 15307 unnumbered argument conversion specifications (that is, % and \*), but not both. The only |  
 15308 exception to this is that %% can be mixed with the "%n\$" form. The results of mixing numbered |  
 15309 and unnumbered argument specifications in a *format* wide-character string are undefined. When |  
 15310 numbered argument specifications are used, specifying the *N*th argument requires that all the |  
 15311 leading arguments, from the first to the (*N*-1)th, are specified in the format wide-character |  
 15312 string.

15313 In *format* wide-character strings containing the "%n\$" form of conversion specification, |  
 15314 numbered arguments in the argument list can be referenced from the *format* wide-character |  
 15315 string as many times as required.

15316 In *format* wide-character strings containing the % form of conversion specification, each |  
 15317 argument in the argument list shall be used exactly once.

15318 CX All forms of the *fwprintf()* function allow for the insertion of a locale-dependent radix character |  
 15319 in the output string, output as a wide-character value. The radix character is defined in the |  
 15320 program's locale (category *LC\_NUMERIC*). In the POSIX locale, or in a locale where the radix |  
 15321 character is not defined, the radix character shall default to a period ('.').

15322 XSI Each conversion specification is introduced by the '%' wide character or by the wide-character  
15323 sequence "%n\$", after which the following appear in sequence:

- 15324 • Zero or more *flags* (in any order), which modify the meaning of the conversion specification.
- 15325 • An optional minimum *field width*. If the converted value has fewer wide characters than the  
15326 field width, it shall be padded with spaces by default on the left; it shall be padded on the  
15327 right, if the left-adjustment flag ('-'), described below, is given to the field width. The field  
15328 width takes the form of an asterisk ('\*'), described below, or a decimal integer.
- 15329 • An optional *precision* that gives the minimum number of digits to appear for the d, i, o, u, x,  
15330 and X conversion specifiers; the number of digits to appear after the radix character for the a,  
15331 A, e, E, f, and F conversion specifiers; the maximum number of significant digits for the g  
15332 and G conversion specifiers; or the maximum number of wide characters to be printed from a  
15333 string in the s conversion specifiers. The precision takes the form of a period ('.') followed  
15334 either by an asterisk ('\*'), described below, or an optional decimal digit string, where a null  
15335 digit string is treated as 0. If a precision appears with any other conversion wide character,  
15336 the behavior is undefined.
- 15337 • An optional length modifier that specifies the size of the argument.
- 15338 • A *conversion specifier* wide character that indicates the type of conversion to be applied.

15339 A field width, or precision, or both, may be indicated by an asterisk ('\*'). In this case an  
15340 argument of type `int` supplies the field width or precision. Applications shall ensure that  
15341 arguments specifying field width, or precision, or both appear in that order before the argument,  
15342 if any, to be converted. A negative field width is taken as a '-' flag followed by a positive field  
15343 XSI width. A negative precision is taken as if the precision were omitted. In format wide-character  
15344 strings containing the "%n\$" form of a conversion specification, a field width or precision may  
15345 be indicated by the sequence "\*m\$", where *m* is a decimal integer in the range  
15346 [1,{NL\_ARGMAX}] giving the position in the argument list (after the format argument) of an  
15347 integer argument containing the field width or precision, for example:

```
15348 wprintf(L"%1$d:%2$.*3$d:%4$.*3$d\n", hour, min, precision, sec);
```

15349 The flag wide characters and their meanings are:

- 15350 XSI ' The integer portion of the result of a decimal conversion (%i, %d, %u, %f, %F, %g, or %G)  
15351 shall be formatted with thousands' grouping wide characters. For other conversions,  
15352 the behavior is undefined. The numeric grouping wide character is used.
- 15353 - The result of the conversion shall be left-justified within the field. The conversion shall  
15354 be right-justified if this flag is not specified.
- 15355 + The result of a signed conversion shall always begin with a sign ('+' or '-'). The  
15356 conversion shall begin with a sign only when a negative value is converted if this flag is  
15357 not specified.
- 15358 <space> If the first wide character of a signed conversion is not a sign, or if a signed conversion  
15359 results in no wide characters, a <space> shall be prefixed to the result. This means that  
15360 if the <space> and '+' flags both appear, the <space> flag shall be ignored.
- 15361 # Specifies that the value is to be converted to an alternative form. For o conversion, it  
15362 increases the precision (if necessary) to force the first digit of the result to be 0. For x or  
15363 X conversion specifiers, a non-zero result shall have 0x (or 0X) prefixed to it. For a, A, e,  
15364 E, f, F, g, and G conversion specifiers, the result shall always contain a radix character,  
15365 even if no digits follow it. Without this flag, a radix character appears in the result of  
15366 these conversions only if a digit follows it. For g and G conversion specifiers, trailing  
15367 zeros shall *not* be removed from the result as they normally are. For other conversion

15368 specifiers, the behavior is undefined. |

15369 0 For `d`, `i`, `o`, `u`, `x`, `X`, `a`, `A`, `e`, `E`, `f`, `F`, `g`, and `G` conversion specifiers, leading zeros |  
15370 (following any indication of sign or base) are used to pad to the field width; no space |  
15371 padding is performed. If the `'0'` and `'-'` flags both appear, the `'0'` flag shall be |  
15372 ignored. For `d`, `i`, `o`, `u`, `x`, and `X` conversion specifiers, if a precision is specified, the `'0'` |  
15373 flag shall be ignored. If the `'0'` and `'\''` flags both appear, the grouping wide |  
15374 characters are inserted before zero padding. For other conversions, the behavior is |  
15375 undefined.

15376 The length modifiers and their meanings are:

15377 `hh` Specifies that a following `d`, `i`, `o`, `u`, `x`, or `X` conversion specifier applies to a **signed char** |  
15378 or **unsigned char** argument (the argument will have been promoted according to the |  
15379 integer promotions, but its value shall be converted to **signed char** or **unsigned char** |  
15380 before printing); or that a following `n` conversion specifier applies to a pointer to a |  
15381 **signed char** argument.

15382 `h` Specifies that a following `d`, `i`, `o`, `u`, `x`, or `X` conversion specifier applies to a **short** or |  
15383 **unsigned short** argument (the argument will have been promoted according to the |  
15384 integer promotions, but its value shall be converted to **short** or **unsigned short** before |  
15385 printing); or that a following `n` conversion specifier applies to a pointer to a **short** |  
15386 argument.

15387 `l` (`ell`) Specifies that a following `d`, `i`, `o`, `u`, `x`, or `X` conversion specifier applies to a **long** or |  
15388 **unsigned long** argument; that a following `n` conversion specifier applies to a pointer to |  
15389 a **long** argument; that a following `c` conversion specifier applies to a **wint\_t** argument; |  
15390 that a following `s` conversion specifier applies to a pointer to a **wchar\_t** argument; or |  
15391 has no effect on a following `a`, `A`, `e`, `E`, `f`, `F`, `g`, or `G` conversion specifier.

15392 `ll` (`ell-ell`) Specifies that a following `d`, `i`, `o`, `u`, `x`, or `X` conversion specifier applies to a **long long** or |  
15393 **unsigned long long** argument; or that a following `n` conversion specifier applies to a |  
15394 pointer to a **long long** argument.

15396 `j` Specifies that a following `d`, `i`, `o`, `u`, `x`, or `X` conversion specifier applies to an **intmax\_t** |  
15397 or **uintmax\_t** argument; or that a following `n` conversion specifier applies to a pointer |  
15398 to an **intmax\_t** argument.

15399 `z` Specifies that a following `d`, `i`, `o`, `u`, `x`, or `X` conversion specifier applies to a **size\_t** or the |  
15400 corresponding signed integer type argument; or that a following `n` conversion specifier |  
15401 applies to a pointer to a signed integer type corresponding to **size\_t** argument.

15402 `t` Specifies that a following `d`, `i`, `o`, `u`, `x`, or `X` conversion specifier applies to a **ptrdiff\_t** or |  
15403 the corresponding **unsigned** type argument; or that a following `n` conversion specifier |  
15404 applies to a pointer to a **ptrdiff\_t** argument.

15405 `L` Specifies that a following `a`, `A`, `e`, `E`, `f`, `F`, `g`, or `G` conversion specifier applies to a **long** |  
15406 **double** argument.

15407 If a length modifier appears with any conversion specifier other than as specified above, the |  
15408 behavior is undefined.

15409 The conversion specifiers and their meanings are:

15410 `d`, `i` The **int** argument shall be converted to a signed decimal in the style "`[-]ddd`". The |  
15411 precision specifies the minimum number of digits to appear; if the value being |  
15412 converted can be represented in fewer digits, it shall be expanded with leading zeros. |  
15413 The default precision shall be 1. The result of converting zero with an explicit precision |

15414		of zero shall be no wide characters.	
15415	o	The <b>unsigned</b> argument shall be converted to unsigned octal format in the style	
15416		"ddd". The precision specifies the minimum number of digits to appear; if the value	
15417		being converted can be represented in fewer digits, it shall be expanded with leading	
15418		zeros. The default precision shall be 1. The result of converting zero with an explicit	
15419		precision of zero shall be no wide characters.	
15420	u	The <b>unsigned</b> argument shall be converted to unsigned decimal format in the style	
15421		"ddd". The precision specifies the minimum number of digits to appear; if the value	
15422		being converted can be represented in fewer digits, it shall be expanded with leading	
15423		zeros. The default precision shall be 1. The result of converting zero with an explicit	
15424		precision of zero shall be no wide characters.	
15425	x	The <b>unsigned</b> argument shall be converted to unsigned hexadecimal format in the style	
15426		"ddd"; the letters "abcdef" are used. The precision specifies the minimum number	
15427		of digits to appear; if the value being converted can be represented in fewer digits, it	
15428		shall be expanded with leading zeros. The default precision shall be 1. The result of	
15429		converting zero with an explicit precision of zero shall be no wide characters.	
15430	X	Equivalent to the x conversion specifier, except that letters "ABCDEF" are used instead	
15431		of "abcdef".	
15432	f, F	The <b>double</b> argument shall be converted to decimal notation in the style	
15433		"[-]ddd.ddd", where the number of digits after the radix character shall be equal to	
15434		the precision specification. If the precision is missing, it shall be taken as 6; if the	
15435		precision is explicitly zero and no '#' flag is present, no radix character shall appear. If	
15436		a radix character appears, at least one digit shall appear before it. The value shall be	
15437		rounded in an implementation-defined manner to the appropriate number of digits.	
15438		A <b>double</b> argument representing an infinity shall be converted in one of the styles	
15439		"[-]inf" or "[-]infinity"; which style is implementation-defined. A <b>double</b>	
15440		argument representing a NaN shall be converted in one of the styles "[-]nan" or	
15441		"[-]nan( <i>n-char-sequence</i> )"; which style, and the meaning of any <i>n-char-sequence</i> ,	
15442		is implementation-defined. The F conversion specifier produces "INF", "INFINITY",	
15443		or "NaN" instead of "inf", "infinity", or "nan", respectively.	
15444	e, E	The <b>double</b> argument shall be converted in the style "[-]d.dde±dd", where there	
15445		shall be one digit before the radix character (which is non-zero if the argument is non-	
15446		zero) and the number of digits after it shall be equal to the precision; if the precision is	
15447		missing, it shall be taken as 6; if the precision is zero and no '#' flag is present, no	
15448		radix character shall appear. The value shall be rounded in an implementation-defined	
15449		manner to the appropriate number of digits. The E conversion wide character shall	
15450		produce a number with 'E' instead of 'e' introducing the exponent. The exponent	
15451		always shall contain at least two digits. If the value is zero, the exponent shall be zero.	
15452		A <b>double</b> argument representing an infinity or NaN shall be converted in the style of	
15453		an f or F conversion specifier.	
15454	g, G	The <b>double</b> argument shall be converted in the style f or e (or in the style F or E in the	
15455		case of a G conversion specifier), with the precision specifying the number of significant	
15456		digits. If an explicit precision is zero, it shall be taken as 1. The style used depends on	
15457		the value converted; style e (or E) shall be used only if the exponent resulting from	
15458		such a conversion is less than -4 or greater than or equal to the precision. Trailing zeros	
15459		shall be removed from the fractional portion of the result; a radix character shall appear	
15460		only if it is followed by a digit.	

15461		A <b>double</b> argument representing an infinity or NaN shall be converted in the style of
15462		an <code>f</code> or <code>F</code> conversion specifier.
15463	a, A	A <b>double</b> argument representing a floating-point number shall be converted in the
15464		style " <code>[-]0xh.hhhhp±d</code> ", where there shall be one hexadecimal digit (which is non-
15465		zero if the argument is a normalized floating-point number and is otherwise
15466		unspecified) before the decimal-point wide character and the number of hexadecimal
15467		digits after it shall be equal to the precision; if the precision is missing and <code>FLT_RADIX</code>
15468		is a power of 2, then the precision shall be sufficient for an exact representation of the
15469		value; if the precision is missing and <code>FLT_RADIX</code> is not a power of 2, then the precision
15470		shall be sufficient to distinguish values of type <b>double</b> , except that trailing zeros may
15471		be omitted; if the precision is zero and the <code>'#'</code> flag is not specified, no decimal-point
15472		wide character shall appear. The letters " <code>abcdef</code> " are used for a conversion and the
15473		letters " <code>ABCDEF</code> " for A conversion. The A conversion specifier produces a number with
15474		<code>'X'</code> and <code>'P'</code> instead of <code>'x'</code> and <code>'p'</code> . The exponent shall always contain at least one
15475		digit, and only as many more digits as necessary to represent the decimal exponent of
15476		2. If the value is zero, the exponent shall be zero.
15477		A <b>double</b> argument representing an infinity or NaN shall be converted in the style of
15478		an <code>f</code> or <code>F</code> conversion specifier.
15479	c	If no <code>l</code> (ell) qualifier is present, the <b>int</b> argument shall be converted to a wide character
15480		as if by calling the <code>btowc()</code> function and the resulting wide character shall be written.
15481		Otherwise, the <b>wint_t</b> argument shall be converted to <b>wchar_t</b> , and written.
15482	s	If no <code>l</code> (ell) qualifier is present, the application shall ensure that the argument is a
15483		pointer to a character array containing a character sequence beginning in the initial
15484		shift state. Characters from the array shall be converted as if by repeated calls to the
15485		<code>mbrtowc()</code> function, with the conversion state described by an <b>mbstate_t</b> object
15486		initialized to zero before the first character is converted, and written up to (but not
15487		including) the terminating null wide character. If the precision is specified, no more
15488		than that many wide characters shall be written. If the precision is not specified, or is
15489		greater than the size of the array, the application shall ensure that the array contains a
15490		null wide character.
15491		If an <code>l</code> (ell) qualifier is present, the application shall ensure that the argument is a
15492		pointer to an array of type <b>wchar_t</b> . Wide characters from the array shall be written up
15493		to (but not including) a terminating null wide character. If no precision is specified, or
15494		is greater than the size of the array, the application shall ensure that the array contains
15495		a null wide character. If a precision is specified, no more than that many wide
15496		characters shall be written.
15497	p	The application shall ensure that the argument is a pointer to <b>void</b> . The value of the
15498		pointer shall be converted to a sequence of printable wide characters in an
15499		implementation-defined manner.
15500	n	The application shall ensure that the argument is a pointer to an integer into which is
15501		written the number of wide characters written to the output so far by this call to one of
15502		the <code>fwprintf()</code> functions. No argument shall be converted, but one shall be consumed. If
15503		the conversion specification includes any flags, a field width, or a precision, the
15504		behavior is undefined.
15505 XSI	C	Equivalent to <code>lc</code> .
15506 XSI	S	Equivalent to <code>ls</code> .

- 15507       %       Output a '%' wide character; no argument shall be converted. The entire conversion |  
 15508                   specification shall be %%. |
- 15509       If a conversion specification does not match one of the above forms, the behavior is undefined.
- 15510       In no case does a nonexistent or small field width cause truncation of a field; if the result of a |  
 15511       conversion is wider than the field width, the field shall be expanded to contain the conversion |  
 15512       result. Characters generated by *fwprintf()* and *wprintf()* shall be printed as if *fputwc()* had been |  
 15513       called.
- 15514       For *a* and *A* conversions, if *FLT\_RADIX* is not a power of 2 and the result is not exactly |  
 15515       representable in the given precision, the result should be one of the two adjacent numbers in |  
 15516       hexadecimal floating style with the given precision, with the extra stipulation that the error |  
 15517       should have a correct sign for the current rounding direction.
- 15518       For *e*, *E*, *f*, *F*, *g*, and *G* conversion specifiers, if the number of significant decimal digits is at most |  
 15519       *DECIMAL\_DIG*, then the result should be correctly rounded. If the number of significant |  
 15520       decimal digits is more than *DECIMAL\_DIG* but the source value is exactly representable with |  
 15521       *DECIMAL\_DIG* digits, then the result should be an exact representation with trailing zeros. |  
 15522       Otherwise, the source value is bounded by two adjacent decimal strings  $L < U$ , both having |  
 15523       *DECIMAL\_DIG* significant digits; the value of the resultant decimal string  $D$  should satisfy  $L \leq$  |  
 15524        $D \leq U$ , with the extra stipulation that the error should have a correct sign for the current |  
 15525       rounding direction.
- 15526 CX       The *st\_ctime* and *st\_mtime* fields of the file shall be marked for update between the call to a |  
 15527       successful execution of *fwprintf()* or *wprintf()* and the next successful completion of a call to |  
 15528       *fflush()* or *fclose()* on the same stream, or a call to *exit()* or *abort()*.
- 15529 **RETURN VALUE**
- 15530       Upon successful completion, these functions shall return the number of wide characters |  
 15531       transmitted, excluding the terminating null wide character in the case of *swprintf()*, or a negative |  
 15532 CX       value if an output error was encountered, and set *errno* to indicate the error.
- 15533       If *n* or more wide characters were requested to be written, *swprintf()* shall return a negative |  
 15534 CX       value, and set *errno* to indicate the error.
- 15535 **ERRORS**
- 15536       For the conditions under which *fwprintf()* and *wprintf()* fail and may fail, refer to *fputwc()*.
- 15537       In addition, all forms of *fwprintf()* may fail if:
- 15538 XSI       [EILSEQ]       A wide-character code that does not correspond to a valid character has been |  
 15539                   detected.
- 15540 XSI       [EINVAL]       There are insufficient arguments.
- 15541       In addition, *wprintf()* and *fwprintf()* may fail if:
- 15542 XSI       [ENOMEM]       Insufficient storage space is available.



15543 **EXAMPLES**

15544 To print the language-independent date and time format, the following statement could be used:

```
15545 wprintf(format, weekday, month, day, hour, min);
```

15546 For American usage, *format* could be a pointer to the wide-character string:

```
15547 L"%s, %s %d, %d:%.2d\n"
```

15548 producing the message:

```
15549 Sunday, July 3, 10:02
```

15550 whereas for German usage, *format* could be a pointer to the wide-character string:

```
15551 L"%1$s, %3$d. %2$s, %4$d:%5$.2d\n"
```

15552 producing the message:

```
15553 Sonntag, 3. Juli, 10:02
```

15554 **APPLICATION USAGE**

15555 None.

15556 **RATIONALE**

15557 None.

15558 **FUTURE DIRECTIONS**

15559 None.

15560 **SEE ALSO**

15561 *btowc()*, *fputwc()*, *fwscanf()*, *mbrtowc()*, *setlocale()*, the Base Definitions volume of  
 15562 IEEE Std 1003.1-200x, `<stdio.h>`, `<wchar.h>`, the Base Definitions volume of  
 15563 IEEE Std 1003.1-200x, Chapter 7, Locale

15564 **CHANGE HISTORY**

15565 First released in Issue 5. Included for alignment with ISO/IEC 9899:1990/Amendment 1:1995  
 15566 (E).

15567 **Issue 6**

15568 The Open Group Corrigendum U040/1 is applied to the RETURN VALUE section, describing  
 15569 the case if *n* or more wide characters are requested to be written using *swprintf()*.

15570 The DESCRIPTION is updated to avoid use of the term “must” for application requirements.

15571 The following changes are made for alignment with the ISO/IEC 9899:1999 standard:

- 15572 • The prototypes for *fwprintf()*, *swprintf()*, and *wprintf()* are updated.
- 15573 • The DESCRIPTION is updated. |
- 15574 • The *hh*, *ll*, *j*, *t*, and *z* length modifiers are added. |
- 15575 • The *a*, *A*, and *F* conversion characters are added. |
- 15576 • XSI shading is removed from the description of character string representations of infinity |
- 15577 and NaN floating-point values. |

15578 The DESCRIPTION is updated to use the terms “conversion specifier” and “conversion  
 15579 specification” consistently. |

15580 ISO/IEC 9899:1999 standard, Technical Corrigendum No. 1 is incorporated. |

15581 **NAME**

15582 fwrite — binary output

15583 **SYNOPSIS**

15584 #include &lt;stdio.h&gt;

15585 size\_t fwrite(const void \*restrict ptr, size\_t size, size\_t nitems,  
15586 FILE \*restrict stream);15587 **DESCRIPTION**15588 CX The functionality described on this reference page is aligned with the ISO C standard. Any  
15589 conflict between the requirements described here and the ISO C standard is unintentional. This  
15590 volume of IEEE Std 1003.1-200x defers to the ISO C standard.15591 The *fwrite()* function shall write, from the array pointed to by *ptr*, up to *nitems* elements whose  
15592 size is specified by *size*, to the stream pointed to by *stream*. For each object, *size* calls shall be  
15593 made to the *fputc()* function, taking the values (in order) from an array of **unsigned char** exactly  
15594 overlaying the object. The file-position indicator for the stream (if defined) shall be advanced by  
15595 the number of bytes successfully written. If an error occurs, the resulting value of the file-  
15596 position indicator for the stream is unspecified.15597 CX The *st\_ctime* and *st\_mtime* fields of the file shall be marked for update between the successful  
15598 execution of *fwrite()* and the next successful completion of a call to *fflush()* or *fclose()* on the  
15599 same stream, or a call to *exit()* or *abort()*.15600 **RETURN VALUE**15601 The *fwrite()* function shall return the number of elements successfully written, which may be  
15602 less than *nitems* if a write error is encountered. If *size* or *nitems* is 0, *fwrite()* shall return 0 and the  
15603 state of the stream remains unchanged. Otherwise, if a write error occurs, the error indicator for  
15604 CX the stream shall be set, and *errno* shall be set to indicate the error.15605 **ERRORS**15606 Refer to *fputc()*.15607 **EXAMPLES**

15608 None.

15609 **APPLICATION USAGE**15610 Because of possible differences in element length and byte ordering, files written using *fwrite()*  
15611 are application-dependent, and possibly cannot be read using *fread()* by a different application  
15612 or by the same application on a different processor.15613 **RATIONALE**

15614 None.

15615 **FUTURE DIRECTIONS**

15616 None.

15617 **SEE ALSO**15618 *ferror()*, *fopen()*, *printf()*, *putc()*, *puts()*, *write()*, the Base Definitions volume of  
15619 IEEE Std 1003.1-200x, <stdio.h>15620 **CHANGE HISTORY**

15621 First released in Issue 1. Derived from Issue 1 of the SVID.

15622 **Issue 6**

15623 Extensions beyond the ISO C standard are now marked.

15624 The following changes are made for alignment with the ISO/IEC 9899:1999 standard:

15625

- The *fwrite()* prototype is updated.

15626

- The DESCRIPTION is updated to clarify how the data is written out using *fputc()*.

## 15627 NAME

15628 fwscanf, swscanf, wscanf — convert formatted wide-character input

## 15629 SYNOPSIS

15630 #include &lt;stdio.h&gt;

15631 #include &lt;wchar.h&gt;

15632 int fwscanf(FILE \*restrict stream, const wchar\_t \*restrict format, ... );

15633 int swscanf(const wchar\_t \*restrict ws,

15634 const wchar\_t \*restrict format, ... );

15635 int wscanf(const wchar\_t \*restrict format, ... );

## 15636 DESCRIPTION

15637 CX The functionality described on this reference page is aligned with the ISO C standard. Any  
 15638 conflict between the requirements described here and the ISO C standard is unintentional. This  
 15639 volume of IEEE Std 1003.1-200x defers to the ISO C standard.

15640 The *fwscanf()* function shall read from the named input *stream*. The *wscanf()* function shall read  
 15641 from the standard input stream *stdin*. The *swscanf()* function shall read from the wide-character  
 15642 string *ws*. Each function reads wide characters, interprets them according to a format, and stores  
 15643 the results in its arguments. Each expects, as arguments, a control wide-character string *format*  
 15644 described below, and a set of *pointer* arguments indicating where the converted input should be  
 15645 stored. The result is undefined if there are insufficient arguments for the format. If the format is  
 15646 exhausted while arguments remain, the excess arguments are evaluated but are otherwise  
 15647 ignored.

15648 XSI Conversions can be applied to the *n*th argument after the *format* in the argument list, rather than  
 15649 to the next unused argument. In this case, the conversion specifier wide character % (see below)  
 15650 is replaced by the sequence "%n\$", where *n* is a decimal integer in the range [1,{NL\_ARGMAX}].  
 15651 This feature provides for the definition of format wide-character strings that select arguments in  
 15652 an order appropriate to specific languages. In format wide-character strings containing the  
 15653 "%n\$" form of conversion specifications, it is unspecified whether numbered arguments in the  
 15654 argument list can be referenced from the format wide-character string more than once.

15655 The *format* can contain either form of a conversion specification—that is, % or "%n\$"—but the  
 15656 two forms cannot normally be mixed within a single *format* wide-character string. The only  
 15657 exception to this is that %% or %\* can be mixed with the "%n\$" form. When numbered  
 15658 argument specifications are used, specifying the *N*th argument requires that all the leading  
 15659 arguments, from the first to the (*N*–1)th, are pointers.

15660 CX The *fwscanf()* function in all its forms allows for detection of a language-dependent radix  
 15661 character in the input string, encoded as a wide-character value. The radix character is defined in  
 15662 the program's locale (category *LC\_NUMERIC*). In the POSIX locale, or in a locale where the  
 15663 radix character is not defined, the radix character shall default to a period ( '.' ).

15664 The *format* is a wide-character string composed of zero or more directives. Each directive is  
 15665 composed of one of the following: one or more white-space wide characters (<space>s, <tab>s,  
 15666 <newline>s, <vertical-tab>s, or <form-feed>s); an ordinary wide character (neither '%' nor a  
 15667 white-space character); or a conversion specification. Each conversion specification is introduced  
 15668 XSI by a '%' or the sequence "%n\$" after which the following appear in sequence:

- 15669 • An optional assignment-suppressing character ' \* '.
- 15670 • An optional non-zero decimal integer that specifies the maximum field width.
- 15671 • An optional length modifier that specifies the size of the receiving object.

15672           • A conversion specifier wide character that specifies the type of conversion to be applied. The  
15673           valid conversion specifiers are described below.

15674           The *fwscanf()* functions shall execute each directive of the format in turn. If a directive fails, as  
15675           detailed below, the function shall return. Failures are described as input failures (due to the  
15676           unavailability of input bytes) or matching failures (due to inappropriate input).

15677           A directive composed of one or more white-space wide characters is executed by reading input  
15678           until no more valid input can be read, or up to the first wide character which is not a white-  
15679           space wide character, which remains unread.

15680           A directive that is an ordinary wide character shall be executed as follows. The next wide  
15681           character is read from the input and compared with the wide character that comprises the  
15682           directive; if the comparison shows that they are not equivalent, the directive shall fail, and the  
15683           differing and subsequent wide characters remain unread. Similarly, if end-of-file, an encoding  
15684           error, or a read error prevents a wide character from being read, the directive shall fail.

15685           A directive that is a conversion specification defines a set of matching input sequences, as  
15686           described below for each conversion wide character. A conversion specification is executed in  
15687           the following steps.

15688           Input white-space wide characters (as specified by *isspace()*) shall be skipped, unless the  
15689           conversion specification includes a `[`, `c`, or `n` conversion specifier.

15690           An item shall be read from the input, unless the conversion specification includes an `n`  
15691           conversion specifier wide character. An input item is defined as the longest sequence of input  
15692           wide characters, not exceeding any specified field width, which is an initial subsequence of a  
15693           matching sequence. The first wide character, if any, after the input item shall remain unread. If  
15694           the length of the input item is zero, the execution of the conversion specification shall fail; this  
15695           condition is a matching failure, unless end-of-file, an encoding error, or a read error prevented  
15696           input from the stream, in which case it is an input failure.

15697           Except in the case of a `%` conversion specifier, the input item (or, in the case of a `%n` conversion  
15698           specification, the count of input wide characters) shall be converted to a type appropriate to the  
15699           conversion wide character. If the input item is not a matching sequence, the execution of the  
15700           conversion specification shall fail; this condition is a matching failure. Unless assignment  
15701           suppression was indicated by a `'*'`, the result of the conversion shall be placed in the object  
15702           pointed to by the first argument following the *format* argument that has not already received a  
15703 XSI           conversion result if the conversion specification is introduced by `%`, or in the *n*th argument if  
15704           introduced by the wide-character sequence `"%n$"`. If this object does not have an appropriate  
15705           type, or if the result of the conversion cannot be represented in the space provided, the behavior  
15706           is undefined.

15707           The length modifiers and their meanings are:

15708           hh       Specifies that a following `d`, `i`, `o`, `u`, `x`, `X`, or `n` conversion specifier applies to an  
15709           argument with type pointer to **signed char** or **unsigned char**.

15710           h        Specifies that a following `d`, `i`, `o`, `u`, `x`, `X`, or `n` conversion specifier applies to an  
15711           argument with type pointer to **short** or **unsigned short**.

15712           l (ell) Specifies that a following `d`, `i`, `o`, `u`, `x`, `X`, or `n` conversion specifier applies to an  
15713           argument with type pointer to **long** or **unsigned long**; that a following `a`, `A`, `e`, `E`, `f`, `F`, `g`,  
15714           or `G` conversion specifier applies to an argument with type pointer to **double**; or that a  
15715           following `c`, `s`, or `[` conversion specifier applies to an argument with type pointer to  
15716           **wchar\_t**.

15717	ll (ell-ell)	
15718		Specifies that a following d, i, o, u, x, X, or n conversion specifier applies to an
15719		argument with type pointer to <b>long long</b> or <b>unsigned long long</b> .
15720	j	Specifies that a following d, i, o, u, x, X, or n conversion specifier applies to an
15721		argument with type pointer to <b>intmax_t</b> or <b>uintmax_t</b> .
15722	z	Specifies that a following d, i, o, u, x, X, or n conversion specifier applies to an
15723		argument with type pointer to <b>size_t</b> or the corresponding signed integer type.
15724	t	Specifies that a following d, i, o, u, x, X, or n conversion specifier applies to an
15725		argument with type pointer to <b>ptrdiff_t</b> or the corresponding <b>unsigned</b> type.
15726	L	Specifies that a following a, A, e, E, f, F, g, or G conversion specifier applies to an
15727		argument with type pointer to <b>long double</b> .
15728		If a length modifier appears with any conversion specifier other than as specified above, the
15729		behavior is undefined.
15730		The following conversion specifier wide characters are valid:
15731	d	Matches an optionally signed decimal integer, whose format is the same as expected for
15732		the subject sequence of <i>wcstol()</i> with the value 10 for the <i>base</i> argument. In the absence
15733		of a size modifier, the application shall ensure that the corresponding argument is a
15734		pointer to <b>int</b> .
15735	i	Matches an optionally signed integer, whose format is the same as expected for the
15736		subject sequence of <i>wcstol()</i> with 0 for the <i>base</i> argument. In the absence of a size
15737		modifier, the application shall ensure that the corresponding argument is a pointer to
15738		<b>int</b> .
15739	o	Matches an optionally signed octal integer, whose format is the same as expected for
15740		the subject sequence of <i>wcstoul()</i> with the value 8 for the <i>base</i> argument. In the absence
15741		of a size modifier, the application shall ensure that the corresponding argument is a
15742		pointer to <b>unsigned</b> .
15743	u	Matches an optionally signed decimal integer, whose format is the same as expected for
15744		the subject sequence of <i>wcstoul()</i> with the value 10 for the <i>base</i> argument. In the absence
15745		of a size modifier, the application shall ensure that the corresponding argument is a
15746		pointer to <b>unsigned</b> .
15747	x	Matches an optionally signed hexadecimal integer, whose format is the same as
15748		expected for the subject sequence of <i>wcstoul()</i> with the value 16 for the <i>base</i> argument.
15749		In the absence of a size modifier, the application shall ensure that the corresponding
15750		argument is a pointer to <b>unsigned</b> .
15751	a, e, f, g	
15752		Matches an optionally signed floating-point number, infinity, or NaN whose format is
15753		the same as expected for the subject sequence of <i>wcstod()</i> . In the absence of a size
15754		modifier, the application shall ensure that the corresponding argument is a pointer to
15755		<b>float</b> .
15756		If the <i>fwprintf()</i> family of functions generates character string representations for
15757		infinity and NaN (a symbolic entity encoded in floating-point format) to support
15758		IEEE Std 754-1985, the <i>fwscanf()</i> family of functions shall recognize them as input.
15759	s	Matches a sequence of non white-space wide characters. If no l (ell) qualifier is present,
15760		characters from the input field shall be converted as if by repeated calls to the
15761		<i>wcrtomb()</i> function, with the conversion state described by an <b>mbstate_t</b> object

15762 initialized to zero before the first wide character is converted. The application shall  
 15763 ensure that the corresponding argument is a pointer to a character array large enough  
 15764 to accept the sequence and the terminating null character, which shall be added  
 15765 automatically.

15766 Otherwise, the application shall ensure that the corresponding argument is a pointer to  
 15767 an array of **wchar\_t** large enough to accept the sequence and the terminating null wide  
 15768 character, which shall be added automatically.

15769 [ Matches a non-empty sequence of wide characters from a set of expected wide  
 15770 characters (the *scanset*). If no **l** (ell) qualifier is present, wide characters from the input  
 15771 field shall be converted as if by repeated calls to the *wcrtomb()* function, with the  
 15772 conversion state described by an **mbstate\_t** object initialized to zero before the first  
 15773 wide character is converted. The application shall ensure that the corresponding  
 15774 argument is a pointer to a character array large enough to accept the sequence and the  
 15775 terminating null character, which shall be added automatically.

15776 If an **l** (ell) qualifier is present, the application shall ensure that the corresponding  
 15777 argument is a pointer to an array of **wchar\_t** large enough to accept the sequence and  
 15778 the terminating null wide character, which shall be added automatically.

15779 The conversion specification includes all subsequent wide characters in the *format*  
 15780 string up to and including the matching right square bracket (']'). The wide  
 15781 characters between the square brackets (the *scanlist*) comprise the scanset, unless the  
 15782 wide character after the left square bracket is a circumflex ('^'), in which case the  
 15783 scanset contains all wide characters that do not appear in the scanlist between the  
 15784 circumflex and the right square bracket. If the conversion specification begins with  
 15785 "[ ]" or "[ ^ ]", the right square bracket is included in the scanlist and the next right  
 15786 square bracket is the matching right square bracket that ends the conversion  
 15787 specification; otherwise, the first right square bracket is the one that ends the  
 15788 conversion specification. If a '-' is in the scanlist and is not the first wide character,  
 15789 nor the second where the first wide character is a '^', nor the last wide character, the  
 15790 behavior is implementation-defined.

15791 **c** Matches a sequence of wide characters of exactly the number specified by the field  
 15792 width (1 if no field width is present in the conversion specification).

15793 If no **l** (ell) length modifier is present, characters from the input field shall be converted  
 15794 as if by repeated calls to the *wcrtomb()* function, with the conversion state described by  
 15795 an **mbstate\_t** object initialized to zero before the first wide character is converted. The  
 15796 corresponding argument shall be a pointer to the initial element of a character array  
 15797 large enough to accept the sequence. No null character is added.

15798 If an **l** (ell) length modifier is present, the corresponding argument shall be a pointer to  
 15799 the initial element of an array of **wchar\_t** large enough to accept the sequence. No null  
 15800 wide character is added.

15801 Otherwise, the application shall ensure that the corresponding argument is a pointer to  
 15802 an array of **wchar\_t** large enough to accept the sequence. No null wide character is  
 15803 added.

15804 **p** Matches an implementation-defined set of sequences, which shall be the same as the set  
 15805 of sequences that is produced by the **%p** conversion specification of the corresponding  
 15806 *fwprintf()* functions. The application shall ensure that the corresponding argument is a  
 15807 pointer to a pointer to **void**. The interpretation of the input item is implementation-  
 15808 defined. If the input item is a value converted earlier during the same program  
 15809 execution, the pointer that results shall compare equal to that value; otherwise, the

15810		behavior of the %p conversion is undefined.
15811	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 wide characters read from the input so far by this call to the <i>fwscanf()</i> functions. Execution of a %n conversion specification shall not increment the assignment count returned at the completion of execution of the function. No argument shall be converted, but one shall be consumed. If the conversion specification includes an assignment-suppressing wide character or a field width, the behavior is undefined.
15812		
15813		
15814		
15815		
15816		
15817		
15818	XSI	C Equivalent to <code>lc</code> .
15819	XSI	S Equivalent to <code>ls</code> .
15820	%	Matches a single '%' wide character; no conversion or assignment shall occur. The complete conversion specification shall be %%.
15821		
15822		If a conversion specification is invalid, the behavior is undefined.
15823		The conversion specifiers A, E, F, G, and X are also valid and shall be equivalent to, respectively,
15824		a, e, f, g, and x.
15825		If end-of-file is encountered during input, conversion is terminated. If end-of-file occurs before any wide characters matching the current conversion specification (except for %n) have been read (other than leading white-space, where permitted), execution of the current conversion specification shall terminate with an input failure. Otherwise, unless execution of the current conversion specification is terminated with a matching failure, execution of the following conversion specification (if any) shall be terminated with an input failure.
15826		
15827		
15828		
15829		
15830		
15831		Reaching the end of the string in <i>swscanf()</i> shall be equivalent to encountering end-of-file for <i>fwscanf()</i> .
15832		
15833		If conversion terminates on a conflicting input, the offending input shall be left unread in the input. Any trailing white space (including <newline>) shall be left unread unless matched by a conversion specification. The success of literal matches and suppressed assignments is only directly determinable via the %n conversion specification.
15834		
15835		
15836		
15837	CX	The <i>fwscanf()</i> and <i>wscanf()</i> functions may mark the <i>st_atime</i> field of the file associated with <i>stream</i> for update. The <i>st_atime</i> field shall be marked for update by the first successful execution of <i>fgetc()</i> , <i>fgetwc()</i> , <i>fgets()</i> , <i>fgetws()</i> , <i>fread()</i> , <i>getc()</i> , <i>getwc()</i> , <i>getchar()</i> , <i>getwchar()</i> , <i>gets()</i> , <i>fscanf()</i> , or <i>fwscanf()</i> using <i>stream</i> that returns data not supplied by a prior call to <i>ungetc()</i> .
15838		
15839		
15840		
15841		<b>RETURN VALUE</b>
15842		Upon successful completion, these functions shall return the number of successfully matched and assigned input items; this number can be zero in the event of an early matching failure. If the input ends before the first matching failure or conversion, EOF shall be returned. If a read error occurs the error indicator for the stream is set, EOF shall be returned, and <i>errno</i> shall be set to indicate the error.
15843		
15844		
15845	CX	
15846		
15847		<b>ERRORS</b>
15848		For the conditions under which the <i>fwscanf()</i> functions shall fail and may fail, refer to <i>fgetwc()</i> .
15849		In addition, <i>fwscanf()</i> may fail if:
15850	XSI	[EILSEQ] Input byte sequence does not form a valid character.
15851	XSI	[EINVAL] There are insufficient arguments.



15852 **EXAMPLES**

15853 The call:

```
15854 int i, n; float x; char name[50];
15855 n = wscanf(L"%d%f%s", &i, &x, name);
```

15856 with the input line:

15857 25 54.32E-1 Hamster

15858 assigns to *n* the value 3, to *i* the value 25, to *x* the value 5.432, and *name* contains the string  
15859 "Hamster".

15860 The call:

```
15861 int i; float x; char name[50];
15862 (void) wscanf(L"%2d%f*d %[0123456789]", &i, &x, name);
```

15863 with input:

15864 56789 0123 56a72

15865 assigns 56 to *i*, 789.0 to *x*, skip 0123, and place the string "56\0" in *name*. The next call to  
15866 *getchar()* shall return the character 'a'.

15867 **APPLICATION USAGE**

15868 In format strings containing the '%' form of conversion specifications, each argument in the  
15869 argument list is used exactly once.

15870 **RATIONALE**

15871 None.

15872 **FUTURE DIRECTIONS**

15873 None.

15874 **SEE ALSO**

15875 *getwc()*, *fwprintf()*, *setlocale()*, *wcstod()*, *wcstol()*, *wcstoul()*, *wcrtomb()*, the Base Definitions  
15876 volume of IEEE Std 1003.1-200x, <langinfo.h>, <stdio.h>, <wchar.h>, the Base Definitions  
15877 volume of IEEE Std 1003.1-200x, Chapter 7, Locale

15878 **CHANGE HISTORY**

15879 First released in Issue 5. Included for alignment with ISO/IEC 9899:1990/Amendment 1:1995  
15880 (E).

15881 **Issue 6**

15882 The DESCRIPTION is updated to avoid use of the term “must” for application requirements.

15883 The following changes are made for alignment with the ISO/IEC 9899:1999 standard:

15884 • The prototypes for *fwscanf()* and *swscanf()* are updated.

15885 • The DESCRIPTION is updated. |

15886 • The hh, ll, j, t, and z length modifiers are added. |

15887 • The a, A, and F conversion characters are added. |

15888 The DESCRIPTION is updated to use the terms “conversion specifier” and “conversion  
15889 specification” consistently.

15890 **NAME**

15891       gai\_strerror — address and name information error description

15892 **SYNOPSIS**

15893       #include &lt;netdb.h&gt;

15894       const char \*gai\_strerror(int *ecode*);15895 **DESCRIPTION**15896       The *gai\_strerror()* function shall return a text string describing an error value for the *getaddrinfo()*  
15897       and *getnameinfo()* functions listed in the <netdb.h> header.15898       When the *ecode* argument is one of the following values listed in the <netdb.h> header:

15899           [EAI\_AGAIN]

15900           [EAI\_BADFLAGS]

15901           [EAI\_FAIL]

15902           [EAI\_FAMILY]

15903           [EAI\_MEMORY]

15904           [EAI\_NONAME]

15905           [EAI\_SERVICE]

15906           [EAI\_SOCKTYPE]

15907           [EAI\_SYSTEM]

15908       the function return value shall point to a string describing the error. If the argument is not one  
15909       of those values, the function shall return a pointer to a string whose contents indicate an  
15910       unknown error.15911 **RETURN VALUE**15912       Upon successful completion, *gai\_strerror()* shall return a pointer to an implementation-defined  
15913       string.15914 **ERRORS**

15915       No errors are defined.

15916 **EXAMPLES**

15917       None.

15918 **APPLICATION USAGE**

15919       None.

15920 **RATIONALE**

15921       None.

15922 **FUTURE DIRECTIONS**

15923       None.

15924 **SEE ALSO**15925       *getaddrinfo()*, the Base Definitions volume of IEEE Std 1003.1-200x, <netdb.h>15926 **CHANGE HISTORY**

15927       First released in Issue 6. Derived from the XNS, Issue 5.2 specification.

15928       The Open Group Base Resolution bwg2001-009 is applied, which changes the return type from

15929       **char \*** to **const char \***. This is for coordination with the IPnG Working Group.

15930 **NAME**15931 `gcvt` — convert a floating-point number to a string (**LEGACY**)15932 **SYNOPSIS**15933 XSI `#include <stdlib.h>`15934 `char *gcvt(double value, int ndigit, char *buf);`

15935

15936 **DESCRIPTION**15937 Refer to `ecvt()`.

15938 **NAME**

15939       getaddrinfo — get address information

15940 **SYNOPSIS**

15941       #include <sys/socket.h>

15942       #include <netdb.h>

15943       int getaddrinfo(const char \*restrict *nodename*,

15944                    const char \*restrict *servname*,

15945                    const struct addrinfo \*restrict *hints*,

15946                    struct addrinfo \*\*restrict *res*);

15947 **DESCRIPTION**

15948       Refer to *freeaddrinfo()*.

15949 **NAME**

15950            getc — get a byte from a stream

15951 **SYNOPSIS**

15952            #include &lt;stdio.h&gt;

15953            int getc(FILE \*stream);

15954 **DESCRIPTION**

15955 **CX**        The functionality described on this reference page is aligned with the ISO C standard. Any  
15956 conflict between the requirements described here and the ISO C standard is unintentional. This  
15957 volume of IEEE Std 1003.1-200x defers to the ISO C standard.

15958        The *getc()* function shall be equivalent to *fgetc()*, except that if it is implemented as a macro it  
15959 may evaluate *stream* more than once, so the argument should never be an expression with side  
15960 effects.

15961 **RETURN VALUE**15962            Refer to *fgetc()*.15963 **ERRORS**15964            Refer to *fgetc()*.15965 **EXAMPLES**

15966            None.

15967 **APPLICATION USAGE**

15968        If the integer value returned by *getc()* is stored into a variable of type **char** and then compared  
15969 against the integer constant EOF, the comparison may never succeed, because sign-extension of  
15970 a variable of type **char** on widening to integer is implementation-defined.

15971        Since it may be implemented as a macro, *getc()* may treat incorrectly a *stream* argument with  
15972 side effects. In particular, *getc(\*f++)* does not necessarily work as expected. Therefore, use of this  
15973 function should be preceded by "#undef getc" in such situations; *fgetc()* could also be used.

15974 **RATIONALE**

15975            None.

15976 **FUTURE DIRECTIONS**

15977            None.

15978 **SEE ALSO**15979            *fgetc()*, the Base Definitions volume of IEEE Std 1003.1-200x, <stdio.h>15980 **CHANGE HISTORY**

15981            First released in Issue 1. Derived from Issue 1 of the SVID.

15982 **NAME**

15983        getc\_unlocked, getchar\_unlocked, putc\_unlocked, putchar\_unlocked — stdio with explicit client  
15984        locking

15985 **SYNOPSIS**

```
15986 TSF     #include <stdio.h>

15987     int getc_unlocked(FILE *stream);
15988     int getchar_unlocked(void);
15989     int putc_unlocked(int c, FILE *stream);
15990     int putchar_unlocked(int c);
15991
```

15992 **DESCRIPTION**

15993        Versions of the functions *getc()*, *getchar()*, *putc()*, and *putchar()* respectively named  
15994        *getc\_unlocked()*, *getchar\_unlocked()*, *putc\_unlocked()*, and *putchar\_unlocked()* shall be provided  
15995        which are functionally equivalent to the original versions, with the exception that they are not  
15996        required to be implemented in a thread-safe manner. They may only safely be used within a  
15997        scope protected by *flockfile()* (or *ftrylockfile()*) and *funlockfile()*. These functions may safely be  
15998        used in a multi-threaded program if and only if they are called while the invoking thread owns  
15999        the (FILE\*) object, as is the case after a successful call of the *flockfile()* or *ftrylockfile()* functions.

16000 **RETURN VALUE**

16001        See *getc()*, *getchar()*, *putc()*, and *putchar()*.

16002 **ERRORS**

16003        See *getc()*, *getchar()*, *putc()*, and *putchar()*.

16004 **EXAMPLES**

16005        None.

16006 **APPLICATION USAGE**

16007        Since they may be implemented as macros, *getc\_unlocked()* and *putc\_unlocked()* may treat  
16008        incorrectly a stream argument with side effects. In particular, *getc\_unlocked(\*f++)* and  
16009        *putc\_unlocked(\*f++)* do not necessarily work as expected. Therefore, use of these functions in  
16010        such situations should be preceded by the following statement as appropriate:

```
16011        #undef getc_unlocked
16012        #undef putc_unlocked
```

16013 **RATIONALE**

16014        Some I/O functions are typically implemented as macros for performance reasons (for example,  
16015        *putc()* and *getc()*). For safety, they need to be synchronized, but it is often too expensive to  
16016        synchronize on every character. Nevertheless, it was felt that the safety concerns were more  
16017        important; consequently, the *getc()*, *getchar()*, *putc()*, and *putchar()* functions are required to be  
16018        thread-safe. However, unlocked versions are also provided with names that clearly indicate the  
16019        unsafe nature of their operation but can be used to exploit their higher performance. These  
16020        unlocked versions can be safely used only within explicitly locked program regions, using  
16021        exported locking primitives. In particular, a sequence such as:

```
16022        flockfile(fileptr);
16023        putc_unlocked('1', fileptr);
16024        putc_unlocked('\n', fileptr);
16025        fprintf(fileptr, "Line 2\n");
16026        funlockfile(fileptr);
```

16027        is permissible, and results in the text sequence:

16028 1  
16029 Line 2  
16030 being printed without being interspersed with output from other threads.  
16031 It would be wrong to have the standard names such as *getc()*, *putc()*, and so on, map to the  
16032 “faster, but unsafe” rather than the “slower, but safe” versions. In either case, you would still  
16033 want to inspect all uses of *getc()*, *putc()*, and so on, by hand when converting existing code.  
16034 Choosing the safe bindings as the default, at least, results in correct code and maintains the  
16035 “atomicity at the function” invariant. To do otherwise would introduce gratuitous  
16036 synchronization errors into converted code. Other routines that modify the *stdio* (**FILE** \*)  
16037 structures or buffers are also safely synchronized.  
16038 Note that there is no need for functions of the form *getc\_locked()*, *putc\_locked()*, and so on, since  
16039 this is the functionality of *getc()*, *putc()*, *et al.* It would be inappropriate to use a feature test  
16040 macro to switch a macro definition of *getc()* between *getc\_locked()* and *getc\_unlocked()*, since the  
16041 ISO C standard requires an actual function to exist, a function whose behavior could not be  
16042 changed by the feature test macro. Also, providing both the *xxx\_locked()* and *xxx\_unlocked()*  
16043 forms leads to the confusion of whether the suffix describes the behavior of the function or the  
16044 circumstances under which it should be used.  
16045 Three additional routines, *flockfile()*, *ftrylockfile()*, and *funlockfile()* (which may be macros), are  
16046 provided to allow the user to delineate a sequence of I/O statements that are executed  
16047 synchronously.  
16048 The *ungetc()* function is infrequently called relative to the other functions/macros so no  
16049 unlocked variation is needed.  
16050 **FUTURE DIRECTIONS**  
16051 None.  
16052 **SEE ALSO**  
16053 *getc()*, *getchar()*, *putc()*, *putchar()*, the Base Definitions volume of IEEE Std 1003.1-200x,  
16054 **<stdio.h>**  
16055 **CHANGE HISTORY**  
16056 First released in Issue 5. Included for alignment with the POSIX Threads Extension.  
16057 **Issue 6**  
16058 These functions are marked as part of the Thread-Safe Functions option.  
16059 The Open Group Corrigendum U030/2 is applied, adding APPLICATION USAGE describing  
16060 how applications should be written to avoid the case when the functions are implemented as  
16061 macros.

16062 **NAME**

16063        getchar — get a byte from a stdin stream

16064 **SYNOPSIS**

16065        #include <stdio.h>

16066        int getchar(void);

16067 **DESCRIPTION**

16068 **cx**        The functionality described on this reference page is aligned with the ISO C standard. Any  
16069        conflict between the requirements described here and the ISO C standard is unintentional. This  
16070        volume of IEEE Std 1003.1-200x defers to the ISO C standard.

16071        The *getchar()* function shall be equivalent to *getc(stdin)*.

16072 **RETURN VALUE**

16073        Refer to *fgetc()*.

16074 **ERRORS**

16075        Refer to *fgetc()*.

16076 **EXAMPLES**

16077        None.

16078 **APPLICATION USAGE**

16079        If the integer value returned by *getchar()* is stored into a variable of type **char** and then  
16080        compared against the integer constant EOF, the comparison may never succeed, because sign-  
16081        extension of a variable of type **char** on widening to integer is implementation-defined.

16082 **RATIONALE**

16083        None.

16084 **FUTURE DIRECTIONS**

16085        None.

16086 **SEE ALSO**

16087        *getc()*, the Base Definitions volume of IEEE Std 1003.1-200x, <**stdio.h**>

16088 **CHANGE HISTORY**

16089        First released in Issue 1. Derived from Issue 1 of the SVID.



16090 **NAME**

16091        getchar\_unlocked — stdio with explicit client locking

16092 **SYNOPSIS**

16093 TSF     #include &lt;stdio.h&gt;

16094        int getchar\_unlocked(void);

16095

16096 **DESCRIPTION**16097        Refer to *getc\_unlocked()*.

## 16098 NAME

16099 getcontext, setcontext — get and set current user context

## 16100 SYNOPSIS

```
16101 xsi      #include <ucontext.h>
16102
16102      int getcontext(ucontext_t *ucp);
16103      int setcontext(const ucontext_t *ucp);
16104
```

## 16105 DESCRIPTION

16106 The *getcontext()* function shall initialize the structure pointed to by *ucp* to the current user  
 16107 context of the calling thread. The **ucontext\_t** type that *ucp* points to defines the user context and  
 16108 includes the contents of the calling thread's machine registers, the signal mask, and the current  
 16109 execution stack.

16110 The *setcontext()* function shall restore the user context pointed to by *ucp*. A successful call to  
 16111 *setcontext()* shall not return; program execution resumes at the point specified by the *ucp*  
 16112 argument passed to *setcontext()*. The *ucp* argument should be created either by a prior call to  
 16113 *getcontext()* or *makecontext()*, or by being passed as an argument to a signal handler. If the *ucp*  
 16114 argument was created with *getcontext()*, program execution continues as if the corresponding  
 16115 call of *getcontext()* had just returned. If the *ucp* argument was created with *makecontext()*,  
 16116 program execution continues with the function passed to *makecontext()*. When that function  
 16117 returns, the thread shall continue as if after a call to *setcontext()* with the *ucp* argument that was  
 16118 input to *makecontext()*. If the *uc\_link* member of the **ucontext\_t** structure pointed to by the *ucp*  
 16119 argument is equal to 0, then this context is the main context, and the thread shall exit when this  
 16120 context returns. The effects of passing a *ucp* argument obtained from any other source are  
 16121 unspecified.

## 16122 RETURN VALUE

16123 Upon successful completion, *setcontext()* shall not return and *getcontext()* shall return 0;  
 16124 otherwise, a value of -1 shall be returned.

## 16125 ERRORS

16126 No errors are defined.

## 16127 EXAMPLES

16128 Refer to *makecontext()*.

## 16129 APPLICATION USAGE

16130 When a signal handler is executed, the current user context is saved and a new context is  
 16131 created. If the thread leaves the signal handler via *longjmp()*, then it is unspecified whether the  
 16132 context at the time of the corresponding *setjmp()* call is restored and thus whether future calls to  
 16133 *getcontext()* provide an accurate representation of the current context, since the context restored  
 16134 by *longjmp()* does not necessarily contain all the information that *setcontext()* requires. Signal  
 16135 handlers should use *siglongjmp()* or *setcontext()* instead.

16136 Conforming applications should not modify or access the *uc\_mcontext* member of **ucontext\_t**. A  
 16137 conforming application cannot assume that context includes any process-wide static data,  
 16138 possibly including *errno*. Users manipulating contexts should take care to handle these  
 16139 explicitly when required.

16140 Use of contexts to create alternate stacks is not defined by this volume of IEEE Std 1003.1-200x.

16141 **RATIONALE**

16142           None.

16143 **FUTURE DIRECTIONS**

16144           None.

16145 **SEE ALSO**16146           *bsd\_signal()*, *makecontext()*, *setjmp()*, *sigaction()*, *sigaltstack()*, *sigprocmask()*, *sigsetjmp()*, the Base  
16147           Definitions volume of IEEE Std 1003.1-200x, <**ucontext.h**>16148 **CHANGE HISTORY**

16149           First released in Issue 4, Version 2.

16150 **Issue 5**

16151           Moved from X/OPEN UNIX extension to BASE.

16152           The following sentence was removed from the DESCRIPTION: “If the *ucp* argument was passed  
16153           to a signal handler, program execution continues with the program instruction following the  
16154           instruction interrupted by the signal.”

16155 **NAME**

16156        getcwd — get the pathname of the current working directory |

16157 **SYNOPSIS**

16158        #include &lt;unistd.h&gt;

16159        char \*getcwd(char \*buf, size\_t size);

16160 **DESCRIPTION**

16161        The *getcwd()* function shall place an absolute pathname of the current working directory in the |  
16162        array pointed to by *buf*, and return *buf*. The pathname copied to the array shall contain no |  
16163        components that are symbolic links. The *size* argument is the size in bytes of the character array |  
16164        pointed to by the *buf* argument. If *buf* is a null pointer, the behavior of *getcwd()* is unspecified.

16165 **RETURN VALUE**

16166        Upon successful completion, *getcwd()* shall return the *buf* argument. Otherwise, *getcwd()* shall |  
16167        return a null pointer and set *errno* to indicate the error. The contents of the array pointed to by |  
16168        *buf* are then undefined.

16169 **ERRORS**16170        The *getcwd()* function shall fail if:16171        [EINVAL]        The *size* argument is 0.

16172        [ERANGE]        The size argument is greater than 0, but is smaller than the length of the |  
16173        pathname +1. |

16174        The *getcwd()* function may fail if:

16175        [EACCES]        Read or search permission was denied for a component of the pathname. |

16176        [ENOMEM]        Insufficient storage space is available.

16177 **EXAMPLES**16178        **Determining the Absolute Pathname of the Current Working Directory** |

16179        The following example returns a pointer to an array that holds the absolute pathname of the |  
16180        current working directory. The pointer is returned in the *ptr* variable, which points to the *buf* |  
16181        array where the pathname is stored. |

16182        #include &lt;stdlib.h&gt;

16183        #include &lt;unistd.h&gt;

16184        ...

16185        long size;

16186        char \*buf;

16187        char \*ptr;

16188        size = pathconf(".", \_PC\_PATH\_MAX);

16189        if ((buf = (char \*)malloc((size\_t)size)) != NULL)

16190            ptr = getcwd(buf, (size\_t)size);

16191        ...

16192 **APPLICATION USAGE**

16193        None.

16194 **RATIONALE**

16195 Since the maximum pathname length is arbitrary unless {PATH\_MAX} is defined, an application  
 16196 generally cannot supply a *buf* with *size* {{PATH\_MAX}+1}.

16197 Having *getcwd()* take no arguments and instead use the *malloc()* function to produce space for  
 16198 the returned argument was considered. The advantage is that *getcwd()* knows how big the  
 16199 working directory pathname is and can allocate an appropriate amount of space. But the  
 16200 programmer would have to use the *free()* function to free the resulting object, or each use of  
 16201 *getcwd()* would further reduce the available memory. Also, *malloc()* and *free()* are used nowhere  
 16202 else in this volume of IEEE Std 1003.1-200x. Finally, *getcwd()* is taken from the SVID where it has  
 16203 the two arguments used in this volume of IEEE Std 1003.1-200x.

16204 The older function *getwd()* was rejected for use in this context because it had only a buffer  
 16205 argument and no *size* argument, and thus had no way to prevent overwriting the buffer, except  
 16206 to depend on the programmer to provide a large enough buffer.

16207 On some implementations, if *buf* is a null pointer, *getcwd()* may obtain *size* bytes of memory  
 16208 using *malloc()*. In this case, the pointer returned by *getcwd()* may be used as the argument in a  
 16209 subsequent call to *free()*. Invoking *getcwd()* with *buf* as a null pointer is not recommended in  
 16210 conforming applications.

16211 If a program is operating in a directory where some (grand)parent directory does not permit  
 16212 reading, *getcwd()* may fail, as in most implementations it must read the directory to determine  
 16213 the name of the file. This can occur if search, but not read, permission is granted in an  
 16214 intermediate directory, or if the program is placed in that directory by some more privileged  
 16215 process (for example, login). Including the [EACCES] error condition makes the reporting of the  
 16216 error consistent and warns the application writer that *getcwd()* can fail for reasons beyond the  
 16217 control of the application writer or user. Some implementations can avoid this occurrence (for  
 16218 example, by implementing *getcwd()* using *pwd*, where *pwd* is a set-user-root process), thus the  
 16219 error was made optional. Since this volume of IEEE Std 1003.1-200x permits the addition of other  
 16220 errors, this would be a common addition and yet one that applications could not be expected to  
 16221 deal with without this addition.

16222 **FUTURE DIRECTIONS**

16223 None.

16224 **SEE ALSO**

16225 *malloc()*, the Base Definitions volume of IEEE Std 1003.1-200x, <unistd.h>

16226 **CHANGE HISTORY**

16227 First released in Issue 1. Derived from Issue 1 of the SVID.

16228 **Issue 6**

16229 The following new requirements on POSIX implementations derive from alignment with the  
 16230 Single UNIX Specification:

- 16231 • The [ENOMEM] optional error condition is added.

16232 **NAME**

16233 getdate — convert user format date and time

16234 **SYNOPSIS**

```
16235 xSI #include <time.h>
16236 struct tm *getdate(const char *string);
16237
```

16238 **DESCRIPTION**

16239 The *getdate()* function shall convert a string representation of a date or time into a broken-down  
 16240 time.

16241 The external variable or macro *getdate\_err* is used by *getdate()* to return error values.

16242 Templates are used to parse and interpret the input string. The templates are contained in a text  
 16243 file identified by the environment variable *DATEMSK*. The *DATEMSK* variable should be set to  
 16244 indicate the full pathname of the file that contains the templates. The first line in the template  
 16245 that matches the input specification is used for interpretation and conversion into the internal  
 16246 time format.

16247 The following conversion specifications shall be supported:

16248	%%	Equivalent to %.	
16249	%a	Abbreviated weekday name.	
16250	%A	Full weekday name.	
16251	%b	Abbreviated month name.	
16252	%B	Full month name.	
16253	%c	Locale's appropriate date and time representation.	
16254	%C	Century number [00,99]; leading zeros are permitted but not required.	
16255	%d	Day of month [01,31]; the leading 0 is optional.	
16256	%D	Date as %m/%d/%y.	
16257	%e	Equivalent to %d.	
16258	%h	Abbreviated month name.	
16259	%H	Hour [00,23].	
16260	%I	Hour [01,12].	
16261	%m	Month number [01,12].	
16262	%M	Minute [00,59].	
16263	%n	Equivalent to <newline>.	
16264	%p	Locale's equivalent of either AM or PM.	
16265	%r	The locale's appropriate representation of time in AM and PM notation. In the POSIX	
16266		locale, this shall be equivalent to %I:%M:%S %p.	
16267	%R	Time as %H:%M.	
16268	%S	Seconds [00,60]. The range goes to 60 (rather than stopping at 59) to allow positive leap	
16269		seconds to be expressed. Since leap seconds cannot be predicted by any algorithm, leap	
16270		second data must come from some external source.	

16271	%t	Equivalent to <tab>.	
16272	%T	Time as %H:%M:%S.	
16273	%w	Weekday number (Sunday = [0,6]).	
16274	%x	Locale's appropriate date representation.	
16275	%X	Locale's appropriate time representation.	
16276	%y	Year within century. When a century is not otherwise specified, values in the range	
16277		[69,99] shall refer to years 1969 to 1999 inclusive, and values in the range [00,68] shall	
16278		refer to years 2000 to 2068 inclusive.	
16279		<b>Note:</b> It is expected that in a future version of IEEE Std 1003.1-200x the default century	
16280		inferred from a 2-digit year will change. (This would apply to all commands	
16281		accepting a 2-digit year as input.)	
16282	%Y	Year as "ccyy" (for example, 2001).	
16283	%Z	Timezone name or no characters if no timezone exists. If the timezone supplied by %Z is	
16284		not the timezone that <i>getdate()</i> expects, an invalid input specification error shall result.	
16285		The <i>getdate()</i> function calculates an expected timezone based on information supplied	
16286		to the function (such as the hour, day, and month).	
16287		The match between the template and input specification performed by <i>getdate()</i> shall be case-	
16288		insensitive.	
16289		The month and weekday names can consist of any combination of upper and lowercase letters.	
16290		The process can request that the input date or time specification be in a specific language by	
16291		setting the <i>LC_TIME</i> category (see <i>setlocale()</i> ).	
16292		Leading 0s are not necessary for the descriptors that allow leading 0s. However, at most two	
16293		digits are allowed for those descriptors, including leading 0s. Extra whitespace in either the	
16294		template file or in <i>string</i> shall be ignored.	
16295		The results are undefined if the conversion specifications %c, %x, and %X include unsupported	
16296		conversion specifications.	
16297		The following rules apply for converting the input specification into the internal format:	
16298		• If %Z is being scanned, then <i>getdate()</i> shall initialize the broken-down time to be the current	
16299		time in the scanned timezone. Otherwise, it shall initialize the broken-down time based on	
16300		the current local time as if <i>localtime()</i> had been called.	
16301		• If only the weekday is given, the day chosen shall be the day, starting with today and moving	
16302		into the future, which first matches the named day.	
16303		• If only the month (and no year) is given, the month chosen shall be the month, starting with	
16304		the current month and moving into the future, which first matches the named month. The	
16305		first day of the month shall be assumed if no day is given.	
16306		• If no hour, minute, and second are given the current hour, minute, and second shall be	
16307		assumed.	
16308		• If no date is given, the hour chosen shall be the hour, starting with the current hour and	
16309		moving into the future, which first matches the named hour.	
16310		If a conversion specification in the <i>DATMSK</i> file does not correspond to one of the conversion	
16311		specifications above, the behavior is unspecified.	
16312		The <i>getdate()</i> function need not be reentrant. A function that is not required to be reentrant is not	
16313		required to be thread-safe.	

16314 **RETURN VALUE**

16315 Upon successful completion, `getdate()` shall return a pointer to a **struct tm**. Otherwise, it shall  
 16316 return a null pointer and set `getdate_err` to indicate the error.

16317 **ERRORS**

16318 The `getdate()` function shall fail in the following cases, setting `getdate_err` to the value shown in  
 16319 the list below. Any changes to `errno` are unspecified.

- 16320 1. The `DATEMSK` environment variable is null or undefined.
- 16321 2. The template file cannot be opened for reading.
- 16322 3. Failed to get file status information.
- 16323 4. The template file is not a regular file.
- 16324 5. An I/O error is encountered while reading the template file.
- 16325 6. Memory allocation failed (not enough memory available).
- 16326 7. There is no line in the template that matches the input.
- 16327 8. Invalid input specification. For example, February 31; or a time is specified that cannot be  
 16328 represented in a `time_t` (representing the time in seconds since the Epoch).

16329 **EXAMPLES**

- 16330 1. The following example shows the possible contents of a template:

```
16331 %m
16332 %A %B %d, %Y, %H:%M:%S
16333 %A
16334 %B
16335 %m/%d/%y %I %p
16336 %d,%m,%Y %H:%M
16337 at %A the %dst of %B in %Y
16338 run job at %I %p,%B %dnd
16339 %A den %d. %B %Y %H.%M Uhr
```

- 16340 2. The following are examples of valid input specifications for the template in Example 1:

```
16341 getdate("10/1/87 4 PM");
16342 getdate("Friday");
16343 getdate("Friday September 18, 1987, 10:30:30");
16344 getdate("24,9,1986 10:30");
16345 getdate("at monday the 1st of december in 1986");
16346 getdate("run job at 3 PM, december 2nd");
```

16347 If the `LC_TIME` category is set to a German locale that includes *freitag* as a weekday name  
 16348 and *oktober* as a month name, the following would be valid:

```
16349 getdate("freitag den 10. oktober 1986 10.30 Uhr");
```

- 16350 3. The following example shows how local date and time specification can be defined in the  
 16351 template:



16352  
16353  
16354  
16355  
16356  
16357

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

16358  
16359  
16360  
16361

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:

16362  
16363  
16364  
16365  
16366  
16367  
16368  
16369  
16370  
16371  
16372  
16373  
16374  
16375

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

#### 16376 APPLICATION USAGE

16377 Although historical versions of *getdate()* did not require that **<time.h>** declare the external  
16378 variable *getdate\_err*, this volume of IEEE Std 1003.1-200x does require it. The standard  
16379 developers encourage applications to remove declarations of *getdate\_err* and instead incorporate  
16380 the declaration by including **<time.h>**.

16381 Applications should use %Y (4-digit years) in preference to %y (2-digit years).

#### 16382 RATIONALE

16383 In standard locales, the conversion specifications %c, %x, and %X do not include unsupported  
16384 conversion specifiers and so the text regarding results being undefined is not a problem in that  
16385 case.

#### 16386 FUTURE DIRECTIONS

16387 None.

#### 16388 SEE ALSO

16389 *ctime()*, *localtime()*, *setlocale()*, *strftime()*, *times()*, the Base Definitions volume of  
16390 IEEE Std 1003.1-200x, **<time.h>**

#### 16391 CHANGE HISTORY

16392 First released in Issue 4, Version 2.

#### 16393 Issue 5

16394 Moved from X/OPEN UNIX extension to BASE.

16395 The last paragraph of the DESCRIPTION is added.

16396 The %C conversion specification is added, and the exact meaning of the %y conversion  
16397 specification is clarified in the DESCRIPTION.

- 16398 A note indicating that this function need not be reentrant is added to the DESCRIPTION.
- 16399 The %R conversion specifications is changed to follow historical practice.
- 16400 **Issue 6**
- 16401 The DESCRIPTION is updated to refer to “seconds since the Epoch” rather than “seconds since  
16402 00:00:00 UTC (Coordinated Universal Time), January 1 1970” for consistency with other *time*  
16403 functions.
- 16404 The description of %S is updated so that the valid range [00,60] rather than [00,61]. |
- 16405 The DESCRIPTION is updated to refer to conversion specifications instead of field descriptors |  
16406 for consistency with other functions.

16407 **NAME**

16408           getegid — get the effective group ID

16409 **SYNOPSIS**

16410           #include <unistd.h>

16411           gid\_t getegid(void);

16412 **DESCRIPTION**

16413           The *getegid()* function shall return the effective group ID of the calling process.

16414 **RETURN VALUE**

16415           The *getegid()* function shall always be successful and no return value is reserved to indicate an error.

16417 **ERRORS**

16418           No errors are defined.

16419 **EXAMPLES**

16420           None.

16421 **APPLICATION USAGE**

16422           None.

16423 **RATIONALE**

16424           None.

16425 **FUTURE DIRECTIONS**

16426           None.

16427 **SEE ALSO**

16428           *geteuid()*, *getgid()*, *getuid()*, *setegid()*, *seteuid()*, *setgid()*, *setregid()*, *setreuid()*, *setuid()*, the Base  
16429           Definitions volume of IEEE Std 1003.1-200x, <sys/types.h>, <unistd.h>

16430 **CHANGE HISTORY**

16431           First released in Issue 1. Derived from Issue 1 of the SVID.

16432 **Issue 6**

16433           In the SYNOPSIS, the optional include of the <sys/types.h> header is removed.

16434           The following new requirements on POSIX implementations derive from alignment with the  
16435           Single UNIX Specification:

- 16436           • The requirement to include <sys/types.h> has been removed. Although <sys/types.h> was  
16437           required for conforming implementations of previous POSIX specifications, it was not  
16438           required for UNIX applications.

16439 **NAME**

16440            getenv — get value of an environment variable

16441 **SYNOPSIS**

16442            #include <stdlib.h>

16443            char \*getenv(const char \*name);

16444 **DESCRIPTION**

16445 **CX**        The functionality described on this reference page is aligned with the ISO C standard. Any  
16446 conflict between the requirements described here and the ISO C standard is unintentional. This  
16447 volume of IEEE Std 1003.1-200x defers to the ISO C standard.

16448        The *getenv()* function shall search the environment of the calling process (see the Base  
16449 Definitions volume of IEEE Std 1003.1-200x, Chapter 8, Environment Variables) for the  
16450 environment variable *name* if it exists and return a pointer to the value of the environment  
16451 variable. If the specified environment variable cannot be found, a null pointer shall be returned.  
16452 The application shall ensure that it does not modify the string pointed to by the *getenv()*  
16453 function.

16454 **CX**        The string pointed to may be overwritten by a subsequent call to *getenv()*, *setenv()*, or *unsetenv()*,  
16455 but shall not be overwritten by a call to any other function in this volume of  
16456 IEEE Std 1003.1-200x.

16457 **CX**        If the application modifies *environ* or the pointers to which it points, the behavior of *getenv()* is  
16458 undefined.

16459        The *getenv()* function need not be reentrant. A function that is not required to be reentrant is not  
16460 required to be thread-safe.

16461 **RETURN VALUE**

16462        Upon successful completion, *getenv()* shall return a pointer to a string containing the *value* for  
16463 the specified *name*. If the specified name cannot be found in the environment of the calling  
16464 process, a null pointer shall be returned.

16465        The return value from *getenv()* may point to static data which may be overwritten by  
16466 **CX**        subsequent calls to *getenv()*, *setenv()*, or *unsetenv()*.

16467 **XSI**        On XSI-conformant systems, the return value from *getenv()* may point to static data which may  
16468 also be overwritten by subsequent calls to *putenv()*.

16469 **ERRORS**

16470        No errors are defined.

16471 **EXAMPLES**16472        **Getting the Value of an Environment Variable**

16473        The following example gets the value of the *HOME* environment variable.

```
16474        #include <stdlib.h>
16475        ...
16476        const char *name = "HOME";
16477        char *value;
16478        value = getenv(name);
```

16479 **APPLICATION USAGE**

16480 None.

16481 **RATIONALE**

16482 The *clearenv()* function was considered but rejected. The *putenv()* function has now been  
 16483 included for alignment with the Single UNIX Specification.

16484 The *getenv()* function is inherently not reentrant because it returns a value pointing to static  
 16485 data.

16486 Conforming applications are required not to modify *environ* directly, but to use only the  
 16487 functions described here to manipulate the process environment as an abstract object. Thus, the  
 16488 implementation of the environment access functions has complete control over the data  
 16489 structure used to represent the environment (subject to the requirement that *environ* be  
 16490 maintained as a list of strings with embedded equal signs for applications that wish to scan the  
 16491 environment). This constraint allows the implementation to properly manage the memory it  
 16492 allocates, either by using allocated storage for all variables (copying them on the first invocation  
 16493 of *setenv()* or *unsetenv()*), or keeping track of which strings are currently in allocated space and  
 16494 which are not, via a separate table or some other means. This enables the implementation to free  
 16495 any allocated space used by strings (and perhaps the pointers to them) stored in *environ* when  
 16496 *unsetenv()* is called. A C runtime start-up procedure (that which invokes *main()* and perhaps  
 16497 initializes *environ*) can also initialize a flag indicating that none of the environment has yet been  
 16498 copied to allocated storage, or that the separate table has not yet been initialized.

16499 In fact, for higher performance of *getenv()*, the implementation could also maintain a separate  
 16500 copy of the environment in a data structure that could be searched much more quickly (such as  
 16501 an indexed hash table, or a binary tree), and update both it and the linear list at *environ* when  
 16502 *setenv()* or *unsetenv()* is invoked.

16503 Performance of *getenv()* can be important for applications which have large numbers of  
 16504 environment variables. Typically, applications like this use the environment as a resource  
 16505 database of user-configurable parameters. The fact that these variables are in the user's shell  
 16506 environment usually means that any other program that uses environment variables (such as *ls*,  
 16507 which attempts to use *COLUMNS*, or really almost any utility (*LANG*, *LC\_ALL*, and so on) is  
 16508 similarly slowed down by the linear search through the variables.

16509 An implementation that maintains separate data structures, or even one that manages the  
 16510 memory it consumes, is not currently required as it was thought it would reduce consensus  
 16511 among implementors who do not want to change their historical implementations.

16512 The POSIX Threads Extension states that multi-threaded applications must not modify *environ*  
 16513 directly, and that IEEE Std 1003.1-200x is providing functions which such applications can use in  
 16514 the future to manipulate the environment in a thread-safe manner. Thus, moving away from  
 16515 application use of *environ* is desirable from that standpoint as well.

16516 **FUTURE DIRECTIONS**

16517 None.

16518 **SEE ALSO**

16519 *exec*, *putenv()*, *setenv()*, *unsetenv()*, the Base Definitions volume of IEEE Std 1003.1-200x,  
 16520 `<stdlib.h>`, the Base Definitions volume of IEEE Std 1003.1-200x, Chapter 8, Environment  
 16521 Variables

16522 **CHANGE HISTORY**

16523 First released in Issue 1. Derived from Issue 1 of the SVID.

16524 **Issue 5**

16525 Normative text previously in the APPLICATION USAGE section is moved to the RETURN  
16526 VALUE section.

16527 A note indicating that this function need not be reentrant is added to the DESCRIPTION.

16528 **Issue 6**

16529 The following changes were made to align with the IEEE P1003.1a draft standard:

16530 • References added to the new *setenv()* and *unsetenv()* functions.

16531 The DESCRIPTION is updated to avoid use of the term “must” for application requirements.

16532 **NAME**

16533           geteuid — get the effective user ID

16534 **SYNOPSIS**

16535           #include <unistd.h>

16536           uid\_t geteuid(void);

16537 **DESCRIPTION**

16538           The *geteuid()* function shall return the effective user ID of the calling process.

16539 **RETURN VALUE**

16540           The *geteuid()* function shall always be successful and no return value is reserved to indicate an error.

16542 **ERRORS**

16543           No errors are defined.

16544 **EXAMPLES**

16545           None.

16546 **APPLICATION USAGE**

16547           None.

16548 **RATIONALE**

16549           None.

16550 **FUTURE DIRECTIONS**

16551           None.

16552 **SEE ALSO**

16553           *getegid()*, *getgid()*, *getuid()*, *setegid()*, *seteuid()*, *setgid()*, *setregid()*, *setreuid()*, *setuid()*, the Base  
16554           Definitions volume of IEEE Std 1003.1-200x, <sys/types.h>, <unistd.h>

16555 **CHANGE HISTORY**

16556           First released in Issue 1. Derived from Issue 1 of the SVID.

16557 **Issue 6**

16558           In the SYNOPSIS, the optional include of the <sys/types.h> header is removed.

16559           The following new requirements on POSIX implementations derive from alignment with the  
16560           Single UNIX Specification:

- 16561           • The requirement to include <sys/types.h> has been removed. Although <sys/types.h> was  
16562           required for conforming implementations of previous POSIX specifications, it was not  
16563           required for UNIX applications.

16564 **NAME**

16565           getgid — get the real group ID

16566 **SYNOPSIS**

16567           #include <unistd.h>

16568           gid\_t getgid(void);

16569 **DESCRIPTION**

16570           The *getgid()* function shall return the real group ID of the calling process.

16571 **RETURN VALUE**

16572           The *getgid()* function shall always be successful and no return value is reserved to indicate an error.

16574 **ERRORS**

16575           No errors are defined.

16576 **EXAMPLES**

16577           None.

16578 **APPLICATION USAGE**

16579           None.

16580 **RATIONALE**

16581           None.

16582 **FUTURE DIRECTIONS**

16583           None.

16584 **SEE ALSO**

16585           *getegid()*, *geteuid()*, *getuid()*, *setegid()*, *seteuid()*, *setgid()*, *setregid()*, *setreuid()*, *setuid()*, the Base Definitions volume of IEEE Std 1003.1-200x, <sys/types.h>, <unistd.h>

16587 **CHANGE HISTORY**

16588           First released in Issue 1. Derived from Issue 1 of the SVID.

16589 **Issue 6**

16590           In the SYNOPSIS, the optional include of the <sys/types.h> header is removed.

16591           The following new requirements on POSIX implementations derive from alignment with the Single UNIX Specification:

- 16592
- 16593           • The requirement to include <sys/types.h> has been removed. Although <sys/types.h> was
  - 16594           required for conforming implementations of previous POSIX specifications, it was not
  - 16595           required for UNIX applications.



16596 **NAME**

16597           getgrent — get the group database entry

16598 **SYNOPSIS**

16599 xSI       #include &lt;grp.h&gt;

16600           struct group \*getgrent(void);

16601

16602 **DESCRIPTION**16603           Refer to *endgrent()*.

## 16604 NAME

16605 getgrgid, getgrgid\_r — get group database entry for a group ID

## 16606 SYNOPSIS

16607 #include <grp.h>

16608 struct group \*getgrgid(gid\_t gid);

16609 TSF int getgrgid\_r(gid\_t gid, struct group \*grp, char \*buffer,

16610 size\_t bufsize, struct group \*\*result);

16611

## 16612 DESCRIPTION

16613 The *getgrgid()* function shall search the group database for an entry with a matching *gid*.

16614 The *getgrgid()* function need not be reentrant. A function that is not required to be reentrant is  
16615 not required to be thread-safe.

16616 TSF The *getgrgid\_r()* function shall update the **group** structure pointed to by *grp* and store a pointer |  
16617 to that structure at the location pointed to by *result*. The structure shall contain an entry from |  
16618 the group database with a matching *gid*. Storage referenced by the group structure is allocated |  
16619 from the memory provided with the *buffer* parameter, which is *bufsize* bytes in size. The |  
16620 maximum size needed for this buffer can be determined with the `{_SC_GETGR_R_SIZE_MAX}` |  
16621 *sysconf()* parameter. A NULL pointer shall be returned at the location pointed to by *result* on |  
16622 error or if the requested entry is not found.

## 16623 RETURN VALUE

16624 Upon successful completion, *getgrgid()* shall return a pointer to a **struct group** with the structure  
16625 defined in `<grp.h>` with a matching entry if one is found. The *getgrgid()* function shall return a  
16626 null pointer if either the requested entry was not found, or an error occurred. On error, *errno*  
16627 shall be set to indicate the error.

16628 The return value may point to a static area which is overwritten by a subsequent call to  
16629 *getgrent()*, *getgrgid()*, or *getgrnam()*.

16630 TSF If successful, the *getgrgid\_r()* function shall return zero; otherwise, an error number shall be  
16631 returned to indicate the error.

## 16632 ERRORS

16633 The *getgrgid()* and *getgrgid\_r()* functions may fail if:

16634 [EIO] An I/O error has occurred.

16635 [EINTR] A signal was caught during *getgrgid()*.

16636 [EMFILE] `{OPEN_MAX}` file descriptors are currently open in the calling process.

16637 [ENFILE] The maximum allowable number of files is currently open in the system.

16638 TSF The *getgrgid\_r()* function may fail if:

16639 TSF [ERANGE] Insufficient storage was supplied via *buffer* and *bufsize* to contain the data to  
16640 be referenced by the resulting **group** structure.

16641 **EXAMPLES**16642 **Finding an Entry in the Group Database**

16643 The following example uses *getgrgid()* to search the group database for a group ID that was  
 16644 previously stored in a **stat** structure, then prints out the group name if it is found. If the group is  
 16645 not found, the program prints the numeric value of the group for the entry.

```
16646 #include <sys/types.h>
16647 #include <grp.h>
16648 #include <stdio.h>
16649 ...
16650 struct stat statbuf;
16651 struct group *grp;
16652 ...
16653 if ((grp = getgrgid(statbuf.st_gid)) != NULL)
16654     printf(" %-8.8s", grp->gr_name);
16655 else
16656     printf(" %-8d", statbuf.st_gid);
16657 ...
```

16658 **APPLICATION USAGE**

16659 Applications wishing to check for error situations should set *errno* to 0 before calling *getgrgid()*.  
 16660 If *errno* is set on return, an error occurred.

16661 The *getgrgid\_r()* function is thread-safe and shall return values in a user-supplied buffer instead  
 16662 of possibly using a static data area that may be overwritten by each call.

16663 **RATIONALE**

16664 None.

16665 **FUTURE DIRECTIONS**

16666 None.

16667 **SEE ALSO**

16668 *endgrent()*, *getgrnam()*, the Base Definitions volume of IEEE Std 1003.1-200x, **<grp.h>**,  
 16669 **<limits.h>**, **<sys/types.h>**

16670 **CHANGE HISTORY**

16671 First released in Issue 1. Derived from System V Release 2.0.

16672 **Issue 5**

16673 Normative text previously in the APPLICATION USAGE section is moved to the RETURN  
 16674 VALUE section.

16675 The *getgrgid\_r()* function is included for alignment with the POSIX Threads Extension.

16676 A note indicating that the *getgrgid()* function need not be reentrant is added to the  
 16677 DESCRIPTION.

16678 **Issue 6**

16679 The *getgrgid\_r()* function is marked as part of the Thread-Safe Functions option.

16680 The Open Group Corrigendum U028/3 is applied, correcting text in the DESCRIPTION  
 16681 describing matching the *gid*.

16682 In the DESCRIPTION, the note about reentrancy is expanded to cover thread-safety.

16683 In the SYNOPSIS, the optional include of the **<sys/types.h>** header is removed.

- 16684 The following new requirements on POSIX implementations derive from alignment with the  
16685 Single UNIX Specification:
- 16686 • The requirement to include `<sys/types.h>` has been removed. Although `<sys/types.h>` was  
16687 required for conforming implementations of previous POSIX specifications, it was not  
16688 required for UNIX applications.
  - 16689 • In the RETURN VALUE section, the requirement to set *errno* on error is added.
  - 16690 • The [EIO], [EINTR], [EMFILE], and [ENFILE] optional error conditions are added.
- 16691 The APPLICATION USAGE section is updated to include a note on the thread-safe function and  
16692 its avoidance of possibly using a static data area.
- 16693 IEEE PASC Interpretation 1003.1 #116 is applied, changing the description of the size of the  
16694 buffer from *bufsize* characters to bytes.

16695 **NAME**

16696 getgrnam, getgrnam\_r — search group database for a name

16697 **SYNOPSIS**

16698 #include &lt;grp.h&gt;

16699 struct group \*getgrnam(const char \*name);

16700 TSF int getgrnam\_r(const char \*name, struct group \*grp, char \*buffer,

16701 size\_t bufsize, struct group \*\*result);

16702

16703 **DESCRIPTION**16704 The *getgrnam()* function shall search the group database for an entry with a matching *name*.16705 The *getgrnam()* function need not be reentrant. A function that is not required to be reentrant is  
16706 not required to be thread-safe.16707 TSF The *getgrnam\_r()* function shall update the **group** structure pointed to by *grp* and store a pointer |  
16708 to that structure at the location pointed to by *result*. The structure shall contain an entry from |  
16709 the group database with a matching *gid* or *name*. Storage referenced by the group structure is |  
16710 allocated from the memory provided with the *buffer* parameter, which is *bufsize* bytes in size. The  
16711 maximum size needed for this buffer can be determined with the `{_SC_GETGR_R_SIZE_MAX}`  
16712 *sysconf()* parameter. A NULL pointer is returned at the location pointed to by *result* on error or if  
16713 the requested entry is not found.16714 **RETURN VALUE**16715 The *getgrnam()* function shall return a pointer to a **struct group** with the structure defined in  
16716 <grp.h> with a matching entry if one is found. The *getgrnam()* function shall return a null  
16717 pointer if either the requested entry was not found, or an error occurred. On error, *errno* shall be  
16718 set to indicate the error.16719 The return value may point to a static area which is overwritten by a subsequent call to  
16720 *getgrent()*, *getgrgid()*, or *getgrnam()*.16721 TSF If successful, the *getgrnam\_r()* function shall return zero; otherwise, an error number shall be  
16722 returned to indicate the error.16723 **ERRORS**16724 The *getgrnam()* and *getgrnam\_r()* functions may fail if:

16725 [EIO] An I/O error has occurred.

16726 [EINTR] A signal was caught during *getgrnam()*.

16727 [EMFILE] {OPEN\_MAX} file descriptors are currently open in the calling process.

16728 [ENFILE] The maximum allowable number of files is currently open in the system.

16729 The *getgrnam\_r()* function may fail if:16730 TSF [ERANGE] Insufficient storage was supplied via *buffer* and *bufsize* to contain the data to  
16731 be referenced by the resulting **group** structure.

16732 **EXAMPLES**

16733 None.

16734 **APPLICATION USAGE**

16735 Applications wishing to check for error situations should set *errno* to 0 before calling *getgrnam()*.  
16736 If *errno* is set on return, an error occurred.

16737 The *getgrnam\_r()* function is thread-safe and shall return values in a user-supplied buffer instead  
16738 of possibly using a static data area that may be overwritten by each call.

16739 **RATIONALE**

16740 None.

16741 **FUTURE DIRECTIONS**

16742 None.

16743 **SEE ALSO**

16744 *endgrent()*, *getgrgid()*, the Base Definitions volume of IEEE Std 1003.1-200x, **<grp.h>**, **<limits.h>**,  
16745 **<sys/types.h>**

16746 **CHANGE HISTORY**

16747 First released in Issue 1. Derived from System V Release 2.0.

16748 **Issue 5**

16749 Normative text previously in the APPLICATION USAGE section is moved to the RETURN  
16750 VALUE section.

16751 The *getgrnam\_r()* function is included for alignment with the POSIX Threads Extension.

16752 A note indicating that the *getgrnam()* function need not be reentrant is added to the  
16753 DESCRIPTION.

16754 **Issue 6**

16755 The *getgrnam\_r()* function is marked as part of the Thread-Safe Functions option.

16756 In the DESCRIPTION, the note about reentrancy is expanded to cover thread-safety.

16757 In the SYNOPSIS, the optional include of the **<sys/types.h>** header is removed.

16758 The following new requirements on POSIX implementations derive from alignment with the  
16759 Single UNIX Specification:

16760 • The requirement to include **<sys/types.h>** has been removed. Although **<sys/types.h>** was  
16761 required for conforming implementations of previous POSIX specifications, it was not  
16762 required for UNIX applications.

16763 • In the RETURN VALUE section, the requirement to set *errno* on error is added.

16764 • The [EIO], [EINTR], [EMFILE], and [ENFILE] optional error conditions are added.

16765 The APPLICATION USAGE section is updated to include a note on the thread-safe function and  
16766 its avoidance of possibly using a static data area.

16767 IEEE PASC Interpretation 1003.1 #116 is applied, changing the description of the size of the  
16768 buffer from *bufsize* characters to bytes.

16769 **NAME**16770 `getgroups` — get supplementary group IDs16771 **SYNOPSIS**16772 `#include <unistd.h>`16773 `int getgroups(int gidsetsize, gid_t grouplist[]);`16774 **DESCRIPTION**

16775 The `getgroups()` function shall fill in the array `grouplist` with the current supplementary group |  
 16776 IDs of the calling process. It is implementation-defined whether `getgroups()` also returns the |  
 16777 effective group ID in the `grouplist` array.

16778 The `gidsetsize` argument specifies the number of elements in the array `grouplist`. The actual |  
 16779 number of group IDs stored in the array shall be returned. The values of array entries with |  
 16780 indices greater than or equal to the value returned are undefined.

16781 If `gidsetsize` is 0, `getgroups()` shall return the number of group IDs that it would otherwise return |  
 16782 without modifying the array pointed to by `grouplist`.

16783 If the effective group ID of the process is returned with the supplementary group IDs, the value |  
 16784 returned shall always be greater than or equal to one and less than or equal to the value of |  
 16785 `{NGROUPS_MAX}+1`.

16786 **RETURN VALUE**

16787 Upon successful completion, the number of supplementary group IDs shall be returned. A |  
 16788 return value of `-1` indicates failure and `errno` shall be set to indicate the error.

16789 **ERRORS**16790 The `getgroups()` function shall fail if:

16791 `[EINVAL]` The `gidsetsize` argument is non-zero and less than the number of group IDs |  
 16792 that would have been returned.

16793 **EXAMPLES**16794 **Getting the Supplementary Group IDs of the Calling Process**

16795 The following example places the current supplementary group IDs of the calling process into |  
 16796 the `group` array.

```
16797 #include <sys/types.h>
16798 #include <unistd.h>
16799 ...
16800 gid_t *group;
16801 int nogroups;
16802 long ngroups_max;

16803 ngroups_max = sysconf(_SC_NGROUPS_MAX) + 1;
16804 group = (gid_t *)malloc(ngroups_max * sizeof(gid_t));
16805 ngroups = getgroups(ngroups_max, group);
```

16806 **APPLICATION USAGE**

16807 None.

16808 **RATIONALE**

16809 The related function `setgroups()` is a privileged operation and therefore is not covered by this |  
 16810 volume of IEEE Std 1003.1-200x.

16811 As implied by the definition of supplementary groups, the effective group ID may appear in the  
16812 array returned by *getgroups()* or it may be returned only by *getegid()*. Duplication may exist, but  
16813 the application needs to call *getegid()* to be sure of getting all of the information. Various  
16814 implementation variations and administrative sequences cause the set of groups appearing in  
16815 the result of *getgroups()* to vary in order and as to whether the effective group ID is included,  
16816 even when the set of groups is the same (in the mathematical sense of “set”). (The history of a  
16817 process and its parents could affect the details of result.)

16818 Applications writers should note that {NGROUPS\_MAX} is not necessarily a constant on all  
16819 implementations.

#### 16820 FUTURE DIRECTIONS

16821 None.

#### 16822 SEE ALSO

16823 *getegid()*, *setgid()*, the Base Definitions volume of IEEE Std 1003.1-200x, <sys/types.h>,  
16824 <unistd.h>

#### 16825 CHANGE HISTORY

16826 First released in Issue 3.

16827 Entry included for alignment with the POSIX.1-1988 standard.

#### 16828 Issue 5

16829 Normative text previously in the APPLICATION USAGE section is moved to the  
16830 DESCRIPTION.

#### 16831 Issue 6

16832 In the SYNOPSIS, the optional include of the <sys/types.h> header is removed.

16833 The following new requirements on POSIX implementations derive from alignment with the  
16834 Single UNIX Specification:

16835 • The requirement to include <sys/types.h> has been removed. Although <sys/types.h> was  
16836 required for conforming implementations of previous POSIX specifications, it was not  
16837 required for UNIX applications.

16838 • A return value of 0 is not permitted, because {NGROUPS\_MAX} cannot be 0. This is a FIPS  
16839 requirement.

16840 The following changes were made to align with the IEEE P1003.1a draft standard:

16841 • Explanation added that the effective group ID may be included in the supplementary group  
16842 list.



16843 **NAME**

16844 gethostbyaddr, gethostbyname — network host database functions

16845 **SYNOPSIS**

16846 #include &lt;netdb.h&gt;

```
16847 OB struct hostent *gethostbyaddr(const void *addr, socklen_t len,
16848 int type);
```

```
16849 struct hostent *gethostbyname(const char *name);
```

16850

16851 **DESCRIPTION**

16852 These functions shall retrieve information about hosts. This information is considered to be |  
 16853 stored in a database that can be accessed sequentially or randomly. Implementation of this |  
 16854 database is unspecified. |

16855 **Note:** In many cases it is implemented by the Domain Name System, as documented in RFC 1034, |  
 16856 RFC 1035, and RFC 1886. |

16857 Entries shall be returned in **hostent** structures. |

16858 The *gethostbyaddr()* function shall return an entry containing addresses of address family *type* for |  
 16859 the host with address *addr*. The *len* argument contains the length of the address pointed to by |  
 16860 *addr*. The *gethostbyaddr()* function need not be reentrant. A function that is not required to be |  
 16861 reentrant is not required to be thread-safe. |

16862 The *gethostbyname()* function shall return an entry containing addresses of address family |  
 16863 AF\_INET for the host with name *name*. The *gethostbyname()* function need not be reentrant. A |  
 16864 function that is not required to be reentrant is not required to be thread-safe. |

16865 The *addr* argument of *gethostbyaddr()* shall be an **in\_addr** structure when *type* is AF\_INET. It |  
 16866 contains a binary format (that is, not null-terminated) address in network byte order. The |  
 16867 *gethostbyaddr()* function is not guaranteed to return addresses of address families other than |  
 16868 AF\_INET, even when such addresses exist in the database. |

16869 If *gethostbyaddr()* returns successfully, then the *h\_addrtype* field in the result shall be the same as |  
 16870 the *type* argument that was passed to the function, and the *h\_addr\_list* field shall list a single |  
 16871 address that is a copy of the *addr* argument that was passed to the function. |

16872 The *name* argument of *gethostbyname()* shall be a node name; the behavior of *gethostbyname()* |  
 16873 when passed a numeric address string is unspecified. For IPv4, a numeric address string shall be |  
 16874 in the dotted-decimal notation described in *inet\_addr()*. |

16875 If *name* is not a numeric address string and is an alias for a valid host name, then *gethostbyname()* |  
 16876 shall return information about the host name to which the alias refers, and *name* shall be |  
 16877 included in the list of aliases returned. |

16878 **RETURN VALUE**

16879 Upon successful completion, these functions shall return a pointer to a **hostent** structure if the |  
 16880 requested entry was found, and a null pointer if the end of the database was reached or the |  
 16881 requested entry was not found. |

16882 Upon unsuccessful completion, *gethostbyaddr()* and *gethostbyname()* shall set *h\_errno* to indicate |  
 16883 the error. |

16884 **ERRORS**

16885 These functions shall fail in the following cases. The *gethostbyaddr()* and *gethostbyname()* |  
 16886 functions shall set *h\_errno* to the value shown in the list below. Any changes to *errno* are |  
 16887 unspecified. |

- 16888 [HOST\_NOT\_FOUND]  
16889 No such host is known.
- 16890 [NO\_DATA] The server recognized the request and the name, but no address is available.  
16891 Another type of request to the name server for the domain might return an  
16892 answer.
- 16893 [NO\_RECOVERY]  
16894 An unexpected server failure occurred which cannot be recovered.
- 16895 [TRY\_AGAIN] A temporary and possibly transient error occurred, such as a failure of a  
16896 server to respond.
- 16897 **EXAMPLES**  
16898 None.
- 16899 **APPLICATION USAGE**  
16900 The *gethostbyaddr()* and *gethostbyname()* functions may return pointers to static data, which may  
16901 be overwritten by subsequent calls to any of these functions.
- 16902 The *getaddrinfo()* and *getnameinfo()* functions are preferred over the *gethostbyaddr()* and  
16903 *gethostbyname()* functions.
- 16904 **RATIONALE**  
16905 None.
- 16906 **FUTURE DIRECTIONS**  
16907 The *gethostbyaddr()* and *gethostbyname()* functions may be withdrawn in a future version.
- 16908 **SEE ALSO**  
16909 *endhostent()*, *endservent()*, *gai\_strerror()*, *getaddrinfo()*, *h\_errno*, *inet\_addr()*, the Base Definitions  
16910 volume of IEEE Std 1003.1-200x, <**netdb.h**>
- 16911 **CHANGE HISTORY**  
16912 First released in Issue 6. Derived from the XNS, Issue 5.2 specification.

16913 **NAME**

16914           gethostbyname — network host database functions

16915 **SYNOPSIS**

16916           #include &lt;netdb.h&gt;

16917 OB        struct hostent \*gethostbyname(const char \*name);

16918

16919 **DESCRIPTION**16920           Refer to *gethostbyaddr()*.

16921 **NAME**

16922       gethostent — network host database functions

16923 **SYNOPSIS**

16924       #include <netdb.h>

16925       struct hostent \*gethostent(void);

16926 **DESCRIPTION**

16927       Refer to *endhostent()*.

16928 **NAME**

16929           gethostid — get an identifier for the current host

16930 **SYNOPSIS**

16931 XSI       #include <unistd.h>

16932           long gethostid(void);

16933

16934 **DESCRIPTION**

16935           The *gethostid()* function shall retrieve a 32-bit identifier for the current host. |

16936 **RETURN VALUE**

16937           Upon successful completion, *gethostid()* shall return an identifier for the current host.

16938 **ERRORS**

16939           No errors are defined.

16940 **EXAMPLES**

16941           None.

16942 **APPLICATION USAGE**

16943           This volume of IEEE Std 1003.1-200x does not define the domain in which the return value is  
16944           unique.

16945 **RATIONALE**

16946           None.

16947 **FUTURE DIRECTIONS**

16948           None.

16949 **SEE ALSO**

16950           *random()*, the Base Definitions volume of IEEE Std 1003.1-200x, <unistd.h>

16951 **CHANGE HISTORY**

16952           First released in Issue 4, Version 2.

16953 **Issue 5**

16954           Moved from X/OPEN UNIX extension to BASE.

16955 **NAME**

16956           gethostname — get name of current host

16957 **SYNOPSIS**

16958           #include &lt;unistd.h&gt;

16959           int gethostname(char \*name, size\_t namelen);

16960 **DESCRIPTION**

16961           The *gethostname()* function shall return the standard host name for the current machine. The  
16962           *namelen* argument shall specify the size of the array pointed to by the *name* argument. The  
16963           returned name shall be null-terminated, except that if *namelen* is an insufficient length to hold  
16964           the host name, then the returned name shall be truncated and it is unspecified whether the  
16965           returned name is null-terminated.

16966           Host names are limited to {HOST\_NAME\_MAX} bytes.

16967 **RETURN VALUE**

16968           Upon successful completion, 0 shall be returned; otherwise, -1 shall be returned.

16969 **ERRORS**

16970           No errors are defined.

16971 **EXAMPLES**

16972           None.

16973 **APPLICATION USAGE**

16974           None.

16975 **RATIONALE**

16976           None.

16977 **FUTURE DIRECTIONS**

16978           None.

16979 **SEE ALSO**16980           *gethostid()*, *uname()*, the Base Definitions volume of IEEE Std 1003.1-200x, <unistd.h>16981 **CHANGE HISTORY**

16982           First released in Issue 6. Derived from the XNS, Issue 5.2 specification.

16983           The Open Group Base Resolution bwg2001-008 is applied, changing the *namelen* parameter from16984           *socklen\_t* to *size\_t*.

16985 **NAME**

16986 getitimer, setitimer — get and set value of interval timer

16987 **SYNOPSIS**

```
16988 XSI #include <sys/time.h>
16989
16989 int getitimer(int which, struct itimerval *value);
16990 int setitimer(int which, const struct itimerval *restrict value,
16991              struct itimerval *restrict ovalue);
16992
```

16993 **DESCRIPTION**

16994 The *getitimer()* function shall store the current value of the timer specified by *which* into the  
 16995 structure pointed to by *value*. The *setitimer()* function shall set the timer specified by *which* to  
 16996 the value specified in the structure pointed to by *value*, and if *ovalue* is not a null pointer, stores  
 16997 the previous value of the timer in the structure pointed to by *ovalue*.

16998 A timer value is defined by the **itimerval** structure, specified in **<sys/time.h>**. If *it\_value* is non-  
 16999 zero, it shall indicate the time to the next timer expiration. If *it\_interval* is non-zero, it shall  
 17000 specify a value to be used in reloading *it\_value* when the timer expires. Setting *it\_value* to 0 shall  
 17001 disable a timer, regardless of the value of *it\_interval*. Setting *it\_interval* to 0 shall disable a timer  
 17002 after its next expiration (assuming *it\_value* is non-zero).

17003 Implementations may place limitations on the granularity of timer values. For each interval  
 17004 timer, if the requested timer value requires a finer granularity than the implementation supports,  
 17005 the actual timer value shall be rounded up to the next supported value.

17006 An XSI-conforming implementation provides each process with at least three interval timers,  
 17007 which are indicated by the *which* argument:

17008	ITIMER_REAL	Decrements in real time. A SIGALRM signal is delivered when this timer expires.
17009		
17010	ITIMER_VIRTUAL	Decrements in process virtual time. It runs only when the process is executing. A SIGVTALRM signal is delivered when it expires.
17011		
17012	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.
17013		
17014		
17015		

17016 The interaction between *setitimer()* and any of *alarm()*, *sleep()*, or *usleep()* is unspecified.

17017 **RETURN VALUE**

17018 Upon successful completion, *getitimer()* or *setitimer()* shall return 0; otherwise,  $-1$  shall be  
 17019 returned and *errno* set to indicate the error.

17020 **ERRORS**

17021 The *setitimer()* function shall fail if:

17022	[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.)
17023		
17024		

17025 The *getitimer()* and *setitimer()* functions may fail if:

17026	[EINVAL]	The <i>which</i> argument is not recognized.
-------	----------	--

17027 **EXAMPLES**

17028           None.

17029 **APPLICATION USAGE**

17030           None.

17031 **RATIONALE**

17032           None.

17033 **FUTURE DIRECTIONS**

17034           None.

17035 **SEE ALSO**

17036           *alarm()*, *sleep()*, *timer\_getoverrun()*, *ualarm()*, *usleep()*, the Base Definitions volume of  
17037           IEEE Std 1003.1-200x, <**signal.h**>, <**sys/time.h**>

17038 **CHANGE HISTORY**

17039           First released in Issue 4, Version 2.

17040 **Issue 5**

17041           Moved from X/OPEN UNIX extension to BASE.

17042 **Issue 6**

17043           The **restrict** keyword is added to the *setitimer()* prototype for alignment with the  
17044           ISO/IEC 9899:1999 standard.



17045 **NAME**

17046           getlogin, getlogin\_r — get login name

17047 **SYNOPSIS**

17048           #include &lt;unistd.h&gt;

17049           char \*getlogin(void);

17050 TSF       int getlogin\_r(char \*name, size\_t namesize);

17051

17052 **DESCRIPTION**

17053       The *getlogin()* function shall return a pointer to a string containing the user name associated by  
 17054       the login activity with the controlling terminal of the current process. If *getlogin()* returns a non-  
 17055       null pointer, then that pointer points to the name that the user logged in under, even if there are  
 17056       several login names with the same user ID.

17057       The *getlogin()* function need not be reentrant. A function that is not required to be reentrant is  
 17058       not required to be thread-safe.

17059 TSF       The *getlogin\_r()* function shall put the name associated by the login activity with the controlling  
 17060       terminal of the current process in the character array pointed to by *name*. The array is *namesize*  
 17061       characters long and should have space for the name and the terminating null character. The  
 17062       maximum size of the login name is {LOGIN\_NAME\_MAX}.

17063       If *getlogin\_r()* is successful, *name* points to the name the user used at login, even if there are  
 17064       several login names with the same user ID.

17065 **RETURN VALUE**

17066       Upon successful completion, *getlogin()* shall return a pointer to the login name or a null pointer  
 17067       if the user's login name cannot be found. Otherwise, it shall return a null pointer and set *errno* to  
 17068       indicate the error.

17069       The return value from *getlogin()* may point to static data whose content is overwritten by each  
 17070       call.

17071 TSF       If successful, the *getlogin\_r()* function shall return zero; otherwise, an error number shall be  
 17072       returned to indicate the error.

17073 **ERRORS**17074       The *getlogin()* and *getlogin\_r()* functions may fail if:

17075       [EMFILE]           {OPEN\_MAX} file descriptors are currently open in the calling process.

17076       [ENFILE]           The maximum allowable number of files is currently open in the system.

17077       [ENXIO]            The calling process has no controlling terminal.

17078       The *getlogin\_r()* function may fail if:

17079 TSF       [ERANGE]        The value of *namesize* is smaller than the length of the string to be returned  
 17080       including the terminating null character.

17081 **EXAMPLES**17082 **Getting the User Login Name**

17083 The following example calls the *getlogin()* function to obtain the name of the user associated  
 17084 with the calling process, and passes this information to the *getpwnam()* function to get the  
 17085 associated user database information.

```

17086 #include <unistd.h>
17087 #include <sys/types.h>
17088 #include <pwd.h>
17089 #include <stdio.h>
17090 ...
17091 char *lgn;
17092 struct passwd *pw;
17093 ...
17094 if ((lgn = getlogin()) == NULL || (pw = getpwnam(lgn)) == NULL) {
17095     fprintf(stderr, "Get of user information failed.\n"); exit(1);
17096 }
```

17097 **APPLICATION USAGE**

17098 Three names associated with the current process can be determined: *getpwuid(geteuid())* shall  
 17099 return the name associated with the effective user ID of the process; *getlogin()* shall return the  
 17100 name associated with the current login activity; and *getpwuid(getuid())* shall return the name  
 17101 associated with the real user ID of the process.

17102 The *getlogin\_r()* function is thread-safe and returns values in a user-supplied buffer instead of  
 17103 possibly using a static data area that may be overwritten by each call.

17104 **RATIONALE**

17105 The *getlogin()* function returns a pointer to the user's login name. The same user ID may be  
 17106 shared by several login names. If it is desired to get the user database entry that is used during  
 17107 login, the result of *getlogin()* should be used to provide the argument to the *getpwnam()*  
 17108 function. (This might be used to determine the user's login shell, particularly where a single user  
 17109 has multiple login shells with distinct login names, but the same user ID.)

17110 The information provided by the *cuserid()* function, which was originally defined in the  
 17111 POSIX.1-1988 standard and subsequently removed, can be obtained by the following:

```
17112 getpwuid(geteuid())
```

17113 while the information provided by historical implementations of *cuserid()* can be obtained by:

```
17114 getpwuid(getuid())
```

17115 The thread-safe version of this function places the user name in a user-supplied buffer and  
 17116 returns a non-zero value if it fails. The non-thread-safe version may return the name in a static  
 17117 data area that may be overwritten by each call.

17118 **FUTURE DIRECTIONS**

17119 None.

17120 **SEE ALSO**

17121 *getpwnam()*, *getpwuid()*, *geteuid()*, *getuid()*, the Base Definitions volume of IEEE Std 1003.1-200x,  
 17122 <limits.h>, <unistd.h>

17123 **CHANGE HISTORY**

17124 First released in Issue 1. Derived from System V Release 2.0.

17125 **Issue 5**

17126 Normative text previously in the APPLICATION USAGE section is moved to the RETURN  
17127 VALUE section.

17128 The *getlogin\_r()* function is included for alignment with the POSIX Threads Extension.

17129 A note indicating that the *getlogin()* function need not be reentrant is added to the  
17130 DESCRIPTION.

17131 **Issue 6**

17132 The *getlogin\_r()* function is marked as part of the Thread-Safe Functions option.

17133 In the DESCRIPTION, the note about reentrancy is expanded to cover thread-safety.

17134 The following new requirements on POSIX implementations derive from alignment with the  
17135 Single UNIX Specification:

17136 • In the RETURN VALUE section, the requirement to set *errno* on error is added.

17137 • The [EMFILE], [ENFILE], and [ENXIO] optional error conditions are added.

17138 The APPLICATION USAGE section is updated to include a note on the thread-safe function and  
17139 its avoidance of possibly using a static data area.

## 17140 NAME

17141 getmsg, getpmsg — receive next message from a STREAMS file (STREAMS)

## 17142 SYNOPSIS

17143 XSR #include &lt;stropts.h&gt;

```

17144 int getmsg(int fildes, struct strbuf *restrict ctlptr,
17145           struct strbuf *restrict dataptr, int *restrict flagsp);
17146 int getpmsg(int fildes, struct strbuf *restrict ctlptr,
17147            struct strbuf *restrict dataptr, int *restrict bandp,
17148            int *restrict flagsp);
17149

```

## 17150 DESCRIPTION

17151 The *getmsg()* function shall retrieve the contents of a message located at the head of the  
 17152 STREAM head read queue associated with a STREAMS file and place the contents into one or  
 17153 more buffers. The message contains either a data part, a control part, or both. The data and  
 17154 control parts of the message shall be placed into separate buffers, as described below. The  
 17155 semantics of each part are defined by the originator of the message.

17156 The *getpmsg()* function shall be equivalent to *getmsg()*, except that it provides finer control over  
 17157 the priority of the messages received. Except where noted, all requirements on *getmsg()* also  
 17158 pertain to *getpmsg()*.

17159 The *fildes* argument specifies a file descriptor referencing a STREAMS-based file.

17160 The *ctlptr* and *dataptr* arguments each point to a **strbuf** structure, in which the *buf* member points  
 17161 to a buffer in which the data or control information is to be placed, and the *maxlen* member  
 17162 indicates the maximum number of bytes this buffer can hold. On return, the *len* member shall  
 17163 contain the number of bytes of data or control information actually received. The *len* member  
 17164 shall be set to 0 if there is a zero-length control or data part and *len* shall be set to -1 if no data or  
 17165 control information is present in the message.

17166 When *getmsg()* is called, *flagsp* should point to an integer that indicates the type of message the  
 17167 process is able to receive. This is described further below.

17168 The *ctlptr* argument is used to hold the control part of the message, and *dataptr* is used to hold  
 17169 the data part of the message. If *ctlptr* (or *dataptr*) is a null pointer or the *maxlen* member is -1, the  
 17170 control (or data) part of the message shall not be processed and shall be left on the STREAM  
 17171 head read queue, and if the *ctlptr* (or *dataptr*) is not a null pointer, *len* shall be set to -1. If the  
 17172 *maxlen* member is set to 0 and there is a zero-length control (or data) part, that zero-length part  
 17173 shall be removed from the read queue and *len* shall be set to 0. If the *maxlen* member is set to 0  
 17174 and there are more than 0 bytes of control (or data) information, that information shall be left on  
 17175 the read queue and *len* shall be set to 0. If the *maxlen* member in *ctlptr* (or *dataptr*) is less than the  
 17176 control (or data) part of the message, *maxlen* bytes shall be retrieved. In this case, the remainder  
 17177 of the message shall be left on the STREAM head read queue and a non-zero return value shall  
 17178 be provided.

17179 By default, *getmsg()* shall process the first available message on the STREAM head read queue.  
 17180 However, a process may choose to retrieve only high-priority messages by setting the integer  
 17181 pointed to by *flagsp* to RS\_HIPRI. In this case, *getmsg()* shall only process the next message if it is  
 17182 a high-priority message. When the integer pointed to by *flagsp* is 0, any available message shall  
 17183 be retrieved. In this case, on return, the integer pointed to by *flagsp* shall be set to RS\_HIPRI if a  
 17184 high-priority message was retrieved, or 0 otherwise.

17185 For *getpmsg()*, the flags are different. The *flagsp* argument points to a bitmask with the following  
 17186 mutually-exclusive flags defined: MSG\_HIPRI, MSG\_BAND, and MSG\_ANY. Like *getmsg()*,

17187 *getpmsg()* shall process the first available message on the STREAM head read queue. A process |  
 17188 may choose to retrieve only high-priority messages by setting the integer pointed to by *flagsp* to |  
 17189 MSG\_HIPRI and the integer pointed to by *bandp* to 0. In this case, *getpmsg()* shall only process |  
 17190 the next message if it is a high-priority message. In a similar manner, a process may choose to |  
 17191 retrieve a message from a particular priority band by setting the integer pointed to by *flagsp* to |  
 17192 MSG\_BAND and the integer pointed to by *bandp* to the priority band of interest. In this case, |  
 17193 *getpmsg()* shall only process the next message if it is in a priority band equal to, or greater than, |  
 17194 the integer pointed to by *bandp*, or if it is a high-priority message. If a process wants to get the |  
 17195 first message off the queue, the integer pointed to by *flagsp* should be set to MSG\_ANY and the |  
 17196 integer pointed to by *bandp* should be set to 0. On return, if the message retrieved was a high- |  
 17197 priority message, the integer pointed to by *flagsp* shall be set to MSG\_HIPRI and the integer |  
 17198 pointed to by *bandp* shall be set to 0. Otherwise, the integer pointed to by *flagsp* shall be set to |  
 17199 MSG\_BAND and the integer pointed to by *bandp* shall be set to the priority band of the message.

17200 If O\_NONBLOCK is not set, *getmsg()* and *getpmsg()* shall block until a message of the type |  
 17201 specified by *flagsp* is available at the front of the STREAM head read queue. If O\_NONBLOCK is |  
 17202 set and a message of the specified type is not present at the front of the read queue, *getmsg()* and |  
 17203 *getpmsg()* shall fail and set *errno* to [EAGAIN].

17204 If a hangup occurs on the STREAM from which messages are retrieved, *getmsg()* and *getpmsg()* |  
 17205 shall continue to operate normally, as described above, until the STREAM head read queue is |  
 17206 empty. Thereafter, they shall return 0 in the *len* members of *ctlptr* and *dataptr*.

#### 17207 RETURN VALUE

17208 Upon successful completion, *getmsg()* and *getpmsg()* shall return a non-negative value. A value |  
 17209 of 0 indicates that a full message was read successfully. A return value of MORECTL indicates |  
 17210 that more control information is waiting for retrieval. A return value of MOREDATA indicates |  
 17211 that more data is waiting for retrieval. A return value of the bitwise-logical OR of MORECTL |  
 17212 and MOREDATA indicates that both types of information remain. Subsequent *getmsg()* and |  
 17213 *getpmsg()* calls shall retrieve the remainder of the message. However, if a message of higher |  
 17214 priority has come in on the STREAM head read queue, the next call to *getmsg()* or *getpmsg()* |  
 17215 shall retrieve that higher-priority message before retrieving the remainder of the previous |  
 17216 message.

17217 If the high priority control part of the message is consumed, the message shall be placed back on |  
 17218 the queue as a normal message of band 0. Subsequent *getmsg()* and *getpmsg()* calls shall retrieve |  
 17219 the remainder of the message. If, however, a priority message arrives or already exists on the |  
 17220 STREAM head, the subsequent call to *getmsg()* or *getpmsg()* shall retrieve the higher-priority |  
 17221 message before retrieving the remainder of the message that was put back.

17222 Upon failure, *getmsg()* and *getpmsg()* shall return -1 and set *errno* to indicate the error.

#### 17223 ERRORS

17224 The *getmsg()* and *getpmsg()* functions shall fail if:

- |       |           |  |
|-------|-----------|--|
| 17225 | [EAGAIN]  | The O_NONBLOCK flag is set and no messages are available.  |
| 17226 | [EBADF]   | The <i>fildes</i> argument is not a valid file descriptor open for reading.  |
| 17227 | [EBADMSG] | The queued message to be read is not valid for <i>getmsg()</i> or <i>getpmsg()</i> or a pending file descriptor is at the STREAM head.                                       |
| 17228 |           |  |
| 17229 | [EINTR]   | A signal was caught during <i>getmsg()</i> or <i>getpmsg()</i> .   |
| 17230 | [EINVAL]  | An illegal value was specified by <i>flagsp</i> , or the STREAM or multiplexer referenced by <i>fildes</i> is linked (directly or indirectly) downstream from a multiplexer. |
| 17231 |           |  |
| 17232 |           |  |

17233 [ENOSTR] A STREAM is not associated with *filde*s.

17234 In addition, *getmsg()* and *getpmsg()* shall fail if the STREAM head had processed an  
 17235 asynchronous error before the call. In this case, the value of *errno* does not reflect the result of  
 17236 *getmsg()* or *getpmsg()* but reflects the prior error.

#### 17237 EXAMPLES

##### 17238 Getting Any Message

17239 In the following example, the value of *fd* is assumed to refer to an open STREAMS file. The call  
 17240 to *getmsg()* retrieves any available message on the associated STREAM-head read queue,  
 17241 returning control and data information to the buffers pointed to by *ctrlbuf* and *databuf*,  
 17242 respectively.

```
17243 #include <stropts.h>
17244 ...
17245 int fd;
17246 char ctrlbuf[128];
17247 char databuf[512];
17248 struct strbuf ctrl;
17249 struct strbuf data;
17250 int flags = 0;
17251 int ret;

17252 ctrl.buf = ctrlbuf;
17253 ctrl.maxlen = sizeof(ctrlbuf);

17254 data.buf = databuf;
17255 data.maxlen = sizeof(databuf);

17256 ret = getmsg (fd, &ctrl, &data, &flags);
```

##### 17257 Getting the First Message off the Queue

17258 In the following example, the call to *getpmsg()* retrieves the first available message on the  
 17259 associated STREAM-head read queue.

```
17260 #include <stropts.h>
17261 ...

17262 int fd;
17263 char ctrlbuf[128];
17264 char databuf[512];
17265 struct strbuf ctrl;
17266 struct strbuf data;
17267 int band = 0;
17268 int flags = MSG_ANY;
17269 int ret;

17270 ctrl.buf = ctrlbuf;
17271 ctrl.maxlen = sizeof(ctrlbuf);

17272 data.buf = databuf;
17273 data.maxlen = sizeof(databuf);

17274 ret = getpmsg (fd, &ctrl, &data, &band, &flags);
```

17275 **APPLICATION USAGE**

17276       None.

17277 **RATIONALE**

17278       None.

17279 **FUTURE DIRECTIONS**

17280       None.

17281 **SEE ALSO**

17282       *poll()*, *putmsg()*, *read()*, *write()*, the Base Definitions volume of IEEE Std 1003.1-200x,  
17283       <**stropts.h**>, Section 2.6 (on page 488)

17284 **CHANGE HISTORY**

17285       First released in Issue 4, Version 2.

17286 **Issue 5**

17287       Moved from X/OPEN UNIX extension to BASE.

17288       A paragraph regarding “high-priority control parts of messages” is added to the RETURN  
17289       VALUE section.

17290 **Issue 6**

17291       This function is marked as part of the XSI STREAMS Option Group.

17292       The **restrict** keyword is added to the *getmsg()* and *getpmsg()* prototypes for alignment with the  
17293       ISO/IEC 9899:1999 standard.

## 17294 NAME

17295 getnameinfo — get name information |

## 17296 SYNOPSIS

17297 #include &lt;sys/socket.h&gt;

17298 #include &lt;netdb.h&gt;

```
17299 int getnameinfo(const struct sockaddr *restrict sa, socklen_t salen,
17300 char *restrict node, socklen_t nodelen, char *restrict service,
17301 socklen_t servicelen, unsigned flags);
```

## 17302 DESCRIPTION

17303 The *getnameinfo()* function shall translate a socket address to a node name and service location, |  
 17304 all of which are defined as in *getaddrinfo()*.

17305 The *sa* argument points to a socket address structure to be translated.

17306 If the *node* argument is non-NULL and the *nodelen* argument is non-zero, then the *node* argument  
 17307 points to a buffer able to contain up to *nodelen* characters that receives the node name as a null-  
 17308 terminated string. If the *node* argument is NULL or the *nodelen* argument is zero, the node name  
 17309 shall not be returned. If the node's name cannot be located, the numeric form of the node's  
 17310 address is returned instead of its name.

17311 If the *service* argument is non-NULL and the *servicelen* argument is non-zero, then the *service*  
 17312 argument points to a buffer able to contain up to *servicelen* bytes that receives the service name |  
 17313 as a null-terminated string. If the *service* argument is NULL or the *servicelen* argument is zero, |  
 17314 the service name shall not be returned. If the service's name cannot be located, the numeric form |  
 17315 of the service address (for example, its port number) shall be returned instead of its name. |

17316 The *flags* argument is a flag that changes the default actions of the function. By default the fully- |  
 17317 qualified domain name (FQDN) for the host shall be returned, but: |

- 17318 • If the flag bit NI\_NOFQDN is set, only the node name portion of the FQDN shall be returned  
 17319 for local hosts.
- 17320 • If the flag bit NI\_NUMERICHOST is set, the numeric form of the host's address shall be  
 17321 returned instead of its name, under all circumstances.
- 17322 • If the flag bit NI\_NAMEREQD is set, an error shall be returned if the host's name cannot be  
 17323 located.
- 17324 • If the flag bit NI\_NUMERICSERV is set, the numeric form of the service address shall be  
 17325 returned (for example, its port number) instead of its name, under all circumstances.
- 17326 • If the flag bit NI\_DGRAM is set, this indicates that the service is a datagram service  
 17327 (SOCK\_DGRAM). The default behavior shall assume that the service is a stream service  
 17328 (SOCK\_STREAM).

## 17329 Notes:

- 17330 1. The two NI\_NUMERICxxx flags are required to support the *-n* flag that many  
 17331 commands provide.
- 17332 2. The NI\_DGRAM flag is required for the few AF\_INET and AF\_INET6 port numbers (for |  
 17333 example, [512,514]) that represent different services for UDP and TCP. |

17334 The *getnameinfo()* function shall be thread-safe.



17335 **RETURN VALUE**

17336 A zero return value for *getnameinfo()* indicates successful completion; a non-zero return value  
 17337 indicates failure. The possible values for the failures are listed in the ERRORS section.

17338 Upon successful completion, *getnameinfo()* shall return the *node* and *service* names, if requested,  
 17339 in the buffers provided. The returned names are always null-terminated strings. |

17340 **ERRORS**

17341 The *getnameinfo()* function shall fail and return the corresponding value if:

17342 [EAI\_AGAIN] The name could not be resolved at this time. Future attempts may succeed.

17343 [EAI\_BADFLAGS]

17344 The *flags* had an invalid value.

17345 [EAI\_FAIL] A non-recoverable error occurred.

17346 [EAI\_FAMILY] The address family was not recognized or the address length was invalid for  
 17347 the specified family.

17348 [EAI\_MEMORY] There was a memory allocation failure.

17349 [EAI\_NONAME] The name does not resolve for the supplied parameters.

17350 NI\_NAMEREQD is set and the host's name cannot be located, or both  
 17351 *nodename* and *servname* were null.

17352 [EAI\_SYSTEM] A system error occurred. The error code can be found in *errno*.

17353 **EXAMPLES**

17354 None.

17355 **APPLICATION USAGE**

17356 If the returned values are to be used as part of any further name resolution (for example, passed  
 17357 to *getaddrinfo()*), applications should provide buffers large enough to store any result possible  
 17358 on the system. |

17359 **RATIONALE**

17360 None.

17361 **FUTURE DIRECTIONS**

17362 None.

17363 **SEE ALSO**

17364 *gai\_strerror()*, *getaddrinfo()*, *getservbyname()*, *getservbyport()*, *inet\_ntop()*, *socket()*, the Base  
 17365 Definitions volume of IEEE Std 1003.1-200x, <netdb.h>, <sys/socket.h>

17366 **CHANGE HISTORY**

17367 First released in Issue 6. Derived from the XNS, Issue 5.2 specification.

17368 The **restrict** keyword is added to the *getnameinfo()* prototype for alignment with the  
 17369 ISO/IEC 9899:1999 standard.

17370 **NAME**

17371           getnetbyaddr — network database functions

17372 **SYNOPSIS**

17373           #include <netdb.h>

17374           struct netent \*getnetbyaddr(uint32\_t net, int type);

17375 **DESCRIPTION**

17376           Refer to *endnetent()*.

17377 **NAME**

17378       getnetbyname — network database functions

17379 **SYNOPSIS**

17380       #include <netdb.h>

17381       struct netent \*getnetbyname(const char \*name);

17382 **DESCRIPTION**

17383       Refer to *endnetent()*.

17384 **NAME**

17385           getnetent — network database functions

17386 **SYNOPSIS**

17387           #include <netdb.h>

17388           struct netent \*getnetent(void);

17389 **DESCRIPTION**

17390           Refer to *endnetent()*.

17391 **NAME**

17392            getopt, optarg, opterr, optind, optopt — command option parsing

17393 **SYNOPSIS**

17394            #include &lt;unistd.h&gt;

17395            int getopt(int argc, char \* const argv[], const char \*optstring);

17396            extern char \*optarg;

17397            extern int optind, opterr, optopt;

17398 **DESCRIPTION**

17399            The *getopt()* function is a command-line parser that shall follow Utility Syntax Guidelines 3, 4, 5, |  
 17400            6, 7, 9, and 10 in the Base Definitions volume of IEEE Std 1003.1-200x, Section 12.2, Utility Syntax |  
 17401            Guidelines. |

17402            The parameters *argc* and *argv* are the argument count and argument array as passed to *main()* |  
 17403            (see *exec*). The argument *optstring* is a string of recognized option characters; if a character is |  
 17404            followed by a colon, the option takes an argument. All option characters allowed by Utility |  
 17405            Syntax Guideline 3 are allowed in *optstring*. The implementation may accept other characters as |  
 17406            an extension.

17407            The variable *optind* is the index of the next element of the *argv[]* vector to be processed. It shall |  
 17408            be initialized to 1 by the system, and *getopt()* shall update it when it finishes with each element |  
 17409            of *argv[]*. When an element of *argv[]* contains multiple option characters, it is unspecified how |  
 17410            *getopt()* determines which options have already been processed.

17411            The *getopt()* function shall return the next option character (if one is found) from *argv* that |  
 17412            matches a character in *optstring*, if there is one that matches. If the option takes an argument, |  
 17413            *getopt()* shall set the variable *optarg* to point to the option-argument as follows:

17414            1. If the option was the last character in the string pointed to by an element of *argv*, then |  
 17415            *optarg* shall contain the next element of *argv*, and *optind* shall be incremented by 2. If the |  
 17416            resulting value of *optind* is greater than *argc*, this indicates a missing option-argument, and |  
 17417            *getopt()* shall return an error indication.

17418            2. Otherwise, *optarg* shall point to the string following the option character in that element of |  
 17419            *argv*, and *optind* shall be incremented by 1. |

17420            If, when *getopt()* is called:

17421            *argv[optind]*    is a null pointer

17422            \**argv[optind]*   is not the character -

17423            *argv[optind]*    points to the string "--"

17424            *getopt()* shall return -1 without changing *optind*. If:

17425            *argv[optind]*    points to the string "--"

17426            *getopt()* shall return -1 after incrementing *optind*.

17427            If *getopt()* encounters an option character that is not contained in *optstring*, it shall return the |  
 17428            question-mark ('?') character. If it detects a missing option-argument, it shall return the colon |  
 17429            character (':') if the first character of *optstring* was a colon, or a question-mark character ('?') |  
 17430            otherwise. In either case, *getopt()* shall set the variable *optopt* to the option character that caused |  
 17431            the error. If the application has not set the variable *opterr* to 0 and the first character of *optstring* |  
 17432            is not a colon, *getopt()* shall also print a diagnostic message to *stderr* in the format specified for |  
 17433            the *getopts* utility.

17434            The *getopt()* function need not be reentrant. A function that is not required to be reentrant is not |  
 17435            required to be thread-safe.

17436 **RETURN VALUE**

17437 The *getopt()* function shall return the next option character specified on the command line.

17438 A colon (':') shall be returned if *getopt()* detects a missing argument and the first character of  
17439 *optstring* was a colon (':').

17440 A question mark ('?') shall be returned if *getopt()* encounters an option character not in  
17441 *optstring* or detects a missing argument and the first character of *optstring* was not a colon (':').

17442 Otherwise, *getopt()* shall return -1 when all command line options are parsed.

17443 **ERRORS**

17444 No errors are defined.

17445 **EXAMPLES**17446 **Parsing Command Line Options**

17447 The following code fragment shows how you might process the arguments for a utility that can  
17448 take the mutually-exclusive options *a* and *b* and the options *f* and *o*, both of which require  
17449 arguments:

```
17450 #include <unistd.h>
17451 int
17452 main(int argc, char *argv[ ])
17453 {
17454     int c;
17455     int bflg, aflag, errflag;
17456     char *ifile;
17457     char *ofile;
17458     extern char *optarg;
17459     extern int optind, optopt;
17460     . . .
17461     while ((c = getopt(argc, argv, ":abf:o:")) != -1) {
17462         switch(c) {
17463             case 'a':
17464                 if (bflg)
17465                     errflag++;
17466                 else
17467                     aflag++;
17468                 break;
17469             case 'b':
17470                 if (aflag)
17471                     errflag++;
17472                 else {
17473                     bflg++;
17474                     bproc();
17475                 }
17476                 break;
17477             case 'f':
17478                 ifile = optarg;
17479                 break;
17480             case 'o':
17481                 ofile = optarg;
17482                 break;
```

```

17483         case ':':          /* -f or -o without operand */
17484             fprintf(stderr,
17485                 "Option -%c requires an operand\n", optopt);
17486             errflg++;
17487             break;
17488         case '?':
17489             fprintf(stderr,
17490                 "Unrecognized option: -%c\n", optopt);
17491             errflg++;
17492     }
17493 }
17494 if (errflg) {
17495     fprintf(stderr, "usage: . . . ");
17496     exit(2);
17497 }
17498 for ( ; optind < argc; optind++) {
17499     if (access(argv[optind], R_OK)) {
17500         . . .
17501     }

```

17502 **This code accepts any of the following as equivalent:**

```

17503 cmd -ao arg path path
17504 cmd -a -o arg path path
17505 cmd -o arg -a path path
17506 cmd -a -o arg -- path path
17507 cmd -a -oarg path path
17508 cmd -aoarg path path

```

### 17509 **Checking Options and Arguments**

17510 The following example parses a set of command line options and prints messages to standard  
17511 output for each option and argument that it encounters.

```

17512 #include <unistd.h>
17513 #include <stdio.h>
17514 ...
17515 int c;
17516 char *filename;
17517 extern char *optarg;
17518 extern int optind, optopt, opterr;
17519 ...
17520 while ((c = getopt(argc, argv, ":abf:")) != -1) {
17521     switch(c) {
17522     case 'a':
17523         printf("a is set\n");
17524         break;
17525     case 'b':
17526         printf("b is set\n");
17527         break;
17528     case 'f':
17529         filename = optarg;
17530         printf("filename is %s\n", filename);
17531         break;

```

```

17532         case ':' :
17533             printf("--%c without filename\n", optopt);
17534             break;
17535         case '?':
17536             printf("unknown arg %c\n", optopt);
17537             break;
17538     }
17539 }

```

#### 17540 **Selecting Options from the Command Line**

17541 The following example selects the type of database routines the user wants to use based on the  
 17542 *Options* argument.

```

17543 #include <unistd.h>
17544 #include <string.h>
17545 ...
17546 char *Options = "hdbt1";
17547 ...
17548 int dbtype, i;
17549 char c;
17550 char *st;
17551 ...
17552 dbtype = 0;
17553 while ((c = getopt(argc, argv, Options)) != -1) {
17554     if ((st = strchr(Options, c)) != NULL) {
17555         dbtype = st - Options;
17556         break;
17557     }
17558 }

```

#### 17559 **APPLICATION USAGE**

17560 The *getopt()* function is only required to support option characters included in Utility Syntax  
 17561 Guideline 3. Many historical implementations of *getopt()* support other characters as options.  
 17562 This is an allowed extension, but applications that use extensions are not maximally portable.  
 17563 Note that support for multi-byte option characters is only possible when such characters can be  
 17564 represented as type **int**.

#### 17565 **RATIONALE**

17566 The *optopt* variable represents historical practice and allows the application to obtain the identity  
 17567 of the invalid option.

17568 The description has been written to make it clear that *getopt()*, like the *getopts* utility, deals with  
 17569 option-arguments whether separated from the option by <blank>s or not. Note that the  
 17570 requirements on *getopt()* and *getopts* are more stringent than the Utility Syntax Guidelines.

17571 The *getopt()* function shall return  $-1$ , rather than EOF, so that <**stdio.h**> is not required.

17572 The special significance of a colon as the first character of *optstring* makes *getopt()* consistent  
 17573 with the *getopts* utility. It allows an application to make a distinction between a missing  
 17574 argument and an incorrect option letter without having to examine the option letter. It is true  
 17575 that a missing argument can only be detected in one case, but that is a case that has to be  
 17576 considered.



17577 **FUTURE DIRECTIONS**

17578 None.

17579 **SEE ALSO**17580 *exec*, the Base Definitions volume of IEEE Std 1003.1-200x, <**unistd.h**>, the Shell and Utilities  
17581 volume of IEEE Std 1003.1-200x17582 **CHANGE HISTORY**

17583 First released in Issue 1. Derived from Issue 1 of the SVID.

17584 **Issue 5**17585 A note indicating that the *getopt()* function need not be reentrant is added to the DESCRIPTION.17586 **Issue 6**

17587 IEEE PASC Interpretation 1003.2 #150 is applied.

17588 **NAME**

17589 getpeername — get the name of the peer socket

17590 **SYNOPSIS**

17591 #include <sys/socket.h>

17592 int getpeername(int *socket*, struct sockaddr \*restrict *address*,  
17593 socklen\_t \*restrict *address\_len*);

17594 **DESCRIPTION**

17595 The *getpeername()* function shall retrieve the peer address of the specified socket, store this  
17596 address in the **sockaddr** structure pointed to by the *address* argument, and store the length of this  
17597 address in the object pointed to by the *address\_len* argument.

17598 If the actual length of the address is greater than the length of the supplied **sockaddr** structure,  
17599 the stored address shall be truncated.

17600 If the protocol permits connections by unbound clients, and the peer is not bound, then the value  
17601 stored in the object pointed to by *address* is unspecified.

17602 **RETURN VALUE**

17603 Upon successful completion, 0 shall be returned. Otherwise, -1 shall be returned and *errno* set to  
17604 indicate the error.

17605 **ERRORS**

17606 The *getpeername()* function shall fail if:

17607 [EBADF] The *socket* argument is not a valid file descriptor.

17608 [EINVAL] The socket has been shut down.

17609 [ENOTCONN] The socket is not connected or otherwise has not had the peer prespecified.

17610 [ENOTSOCK] The *socket* argument does not refer to a socket.

17611 [EOPNOTSUPP] The operation is not supported for the socket protocol.

17612 The *getpeername()* function may fail if:

17613 [ENOBUFS] Insufficient resources were available in the system to complete the call.

17614 **EXAMPLES**

17615 None.

17616 **APPLICATION USAGE**

17617 None.

17618 **RATIONALE**

17619 None.

17620 **FUTURE DIRECTIONS**

17621 None.

17622 **SEE ALSO**

17623 *accept()*, *bind()*, *getsockname()*, *socket()*, the Base Definitions volume of IEEE Std 1003.1-200x,  
17624 <sys/socket.h>

17625 **CHANGE HISTORY**

17626 First released in Issue 6. Derived from the XNS, Issue 5.2 specification.

17627 The **restrict** keyword is added to the *getpeername()* prototype for alignment with the  
17628 ISO/IEC 9899:1999 standard.

17629 **NAME**

17630 getpgid — get the process group ID for a process

17631 **SYNOPSIS**

17632 XSI #include <unistd.h>

17633 pid\_t getpgid(pid\_t pid);

17634

17635 **DESCRIPTION**

17636 The *getpgid()* function shall return the process group ID of the process whose process ID is equal  
17637 to *pid*. If *pid* is equal to 0, *getpgid()* shall return the process group ID of the calling process.

17638 **RETURN VALUE**

17639 Upon successful completion, *getpgid()* shall return a process group ID. Otherwise, it shall return  
17640 (**pid\_t**)-1 and set *errno* to indicate the error.

17641 **ERRORS**

17642 The *getpgid()* function shall fail if:

17643 [EPERM] The process whose process ID is equal to *pid* is not in the same session as the  
17644 calling process, and the implementation does not allow access to the process  
17645 group ID of that process from the calling process.

17646 [ESRCH] There is no process with a process ID equal to *pid*.

17647 The *getpgid()* function may fail if:

17648 [EINVAL] The value of the *pid* argument is invalid.

17649 **EXAMPLES**

17650 None.

17651 **APPLICATION USAGE**

17652 None.

17653 **RATIONALE**

17654 None.

17655 **FUTURE DIRECTIONS**

17656 None.

17657 **SEE ALSO**

17658 *exec*, *fork()*, *getpgrp()*, *getpid()*, *getsid()*, *setpgid()*, *setsid()*, the Base Definitions volume of  
17659 IEEE Std 1003.1-200x, <unistd.h>

17660 **CHANGE HISTORY**

17661 First released in Issue 4, Version 2.

17662 **Issue 5**

17663 Moved from X/OPEN UNIX extension to BASE.

17664 **NAME**

17665           getpgrp — get the process group ID of the calling process

17666 **SYNOPSIS**

17667           #include <unistd.h>  
17668           pid\_t getpgrp(void);

17669 **DESCRIPTION**

17670           The *getpgrp()* function shall return the process group ID of the calling process.

17671 **RETURN VALUE**

17672           The *getpgrp()* function shall always be successful and no return value is reserved to indicate an |  
17673           error. |

17674 **ERRORS**

17675           No errors are defined.

17676 **EXAMPLES**

17677           None.

17678 **APPLICATION USAGE**

17679           None.

17680 **RATIONALE**

17681           4.3 BSD provides a *getpgrp()* function that returns the process group ID for a specified process. |  
17682           Although this function supports job control, all known job control shells always specify the |  
17683           calling process with this function. Thus, the simpler System V *getpgrp()* suffices, and the added |  
17684           complexity of the 4.3 BSD *getpgrp()* is provided by the XSI extension *getpgid()*.

17685 **FUTURE DIRECTIONS**

17686           None.

17687 **SEE ALSO**

17688           *exec*, *fork()*, *getpgid()*, *getpid()*, *getppid()*, *kill()*, *setpgid()*, *setsid()*, the Base Definitions volume of  
17689           IEEE Std 1003.1-200x, <sys/types.h>, <unistd.h>

17690 **CHANGE HISTORY**

17691           First released in Issue 1. Derived from Issue 1 of the SVID.

17692 **Issue 6**

17693           In the SYNOPSIS, the optional include of the <sys/types.h> header is removed.

17694           The following new requirements on POSIX implementations derive from alignment with the  
17695           Single UNIX Specification:

- 17696
  - The requirement to include <sys/types.h> has been removed. Although <sys/types.h> was  
17697           required for conforming implementations of previous POSIX specifications, it was not  
17698           required for UNIX applications.

17699 **NAME**

17700            *getpid* — get the process ID

17701 **SYNOPSIS**

17702            #include <unistd.h>

17703            pid\_t *getpid*(void);

17704 **DESCRIPTION**

17705            The *getpid*() function shall return the process ID of the calling process.

17706 **RETURN VALUE**

17707            The *getpid*() function shall always be successful and no return value is reserved to indicate an error.

17709 **ERRORS**

17710            No errors are defined.

17711 **EXAMPLES**

17712            None.

17713 **APPLICATION USAGE**

17714            None.

17715 **RATIONALE**

17716            None.

17717 **FUTURE DIRECTIONS**

17718            None.

17719 **SEE ALSO**

17720            *exec*, *fork*(), *getpgrp*(), *getppid*(), *kill*(), *setpgid*(), *setsid*(), the Base Definitions volume of IEEE Std 1003.1-200x, <sys/types.h>, <unistd.h>

17722 **CHANGE HISTORY**

17723            First released in Issue 1. Derived from Issue 1 of the SVID.

17724 **Issue 6**

17725            In the SYNOPSIS, the optional include of the <sys/types.h> header is removed.

17726            The following new requirements on POSIX implementations derive from alignment with the Single UNIX Specification:

- 17727
- 17728            • 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.
- 17730

17731 **NAME**

17732       getpmsg — receive next message from a STREAMS file

17733 **SYNOPSIS**

```
17734 xSI     #include <stropts.h>
```

```
17735       int getpmsg(int fildev, struct strbuf *restrict ctlptr,
17736                   struct strbuf *restrict dataptr, int *restrict bandp,
17737                   int *restrict flagsp);
```

17738

17739 **DESCRIPTION**

17740       Refer to *getmsg()*.

17741 **NAME**

17742           getppid — get the parent process ID

17743 **SYNOPSIS**

17744           #include <unistd.h>

17745           pid\_t getppid(void);

17746 **DESCRIPTION**

17747           The *getppid()* function shall return the parent process ID of the calling process.

17748 **RETURN VALUE**

17749           The *getppid()* function shall always be successful and no return value is reserved to indicate an error.

17751 **ERRORS**

17752           No errors are defined.

17753 **EXAMPLES**

17754           None.

17755 **APPLICATION USAGE**

17756           None.

17757 **RATIONALE**

17758           None.

17759 **FUTURE DIRECTIONS**

17760           None.

17761 **SEE ALSO**

17762           *exec*, *fork()*, *getpgid()*, *getpgrp()*, *getpid()*, *kill()*, *setpgid()*, *setsid()*, the Base Definitions volume of IEEE Std 1003.1-200x, <sys/types.h>, <unistd.h>

17764 **CHANGE HISTORY**

17765           First released in Issue 1. Derived from Issue 1 of the SVID.

17766 **Issue 6**

17767           In the SYNOPSIS, the optional include of the <sys/types.h> header is removed.

17768           The following new requirements on POSIX implementations derive from alignment with the Single UNIX Specification:

- 17770
  - 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.
- 17771
- 17772

## 17773 NAME

17774 getpriority, setpriority — get and set the nice value

## 17775 SYNOPSIS

17776 xSI #include &lt;sys/resource.h&gt;

17777 int getpriority(int which, id\_t who);

17778 int setpriority(int which, id\_t who, int value);

17779

## 17780 DESCRIPTION

17781 The *getpriority()* function shall obtain the nice value of a process, process group, or user. The  
 17782 *setpriority()* function shall set the nice value of a process, process group, or user to  
 17783 *value*+{NZERO}.

17784 Target processes are specified by the values of the *which* and *who* arguments. The *which*  
 17785 argument may be one of the following values: PRIO\_PROCESS, PRIO\_PGRP, or PRIO\_USER,  
 17786 indicating that the *who* argument is to be interpreted as a process ID, a process group ID, or an  
 17787 effective user ID, respectively. A 0 value for the *who* argument specifies the current process,  
 17788 process group, or user.

17789 The nice value set with *setpriority()* shall be applied to the process. If the process is multi-  
 17790 threaded, the nice value shall affect all system scope threads in the process.

17791 If more than one process is specified, *getpriority()* shall return value {NZERO} less than the  
 17792 lowest nice value pertaining to any of the specified processes, and *setpriority()* shall set the nice  
 17793 values of all of the specified processes to *value*+{NZERO}.

17794 The default nice value is {NZERO}; lower nice values shall cause more favorable scheduling.  
 17795 While the range of valid nice values is [0,{NZERO}\*2-1], implementations may enforce more  
 17796 restrictive limits. If *value*+{NZERO} is less than the system's lowest supported nice value,  
 17797 *setpriority()* shall set the nice value to the lowest supported value; if *value*+{NZERO} is greater  
 17798 than the system's highest supported nice value, *setpriority()* shall set the nice value to the highest  
 17799 supported value.

17800 Only a process with appropriate privileges can lower its nice value.

17801 PS|TPS Any processes or threads using SCHED\_FIFO or SCHED\_RR shall be unaffected by a call to  
 17802 *setpriority()*. This is not considered an error. A process which subsequently reverts to  
 17803 SCHED\_OTHER need not have its priority affected by such a *setpriority()* call.

17804 The effect of changing the nice value may vary depending on the process-scheduling algorithm  
 17805 in effect.

17806 Since *getpriority()* can return the value -1 on successful completion, it is necessary to set *errno* to  
 17807 0 prior to a call to *getpriority()*. If *getpriority()* returns the value -1, then *errno* can be checked to  
 17808 see if an error occurred or if the value is a legitimate nice value.

## 17809 RETURN VALUE

17810 Upon successful completion, *getpriority()* shall return an integer in the range from -{NZERO} to  
 17811 {NZERO}-1. Otherwise, -1 shall be returned and *errno* set to indicate the error.

17812 Upon successful completion, *setpriority()* shall return 0; otherwise, -1 shall be returned and *errno*  
 17813 set to indicate the error.

## 17814 ERRORS

17815 The *getpriority()* and *setpriority()* functions shall fail if:

17816 [ESRCH] No process could be located using the *which* and *who* argument values  
 17817 specified.



17818 [EINVAL] The value of the *which* argument was not recognized, or the value of the *who*  
 17819 argument is not a valid process ID, process group ID, or user ID.

17820 In addition, *setpriority()* may fail if:

17821 [EPERM] A process was located, but neither the real nor effective user ID of the  
 17822 executing process match the effective user ID of the process whose nice value  
 17823 is being changed.

17824 [EACCES] A request was made to change the nice value to a lower numeric value and  
 17825 the current process does not have appropriate privileges.

#### 17826 EXAMPLES

##### 17827 Using *getpriority()*

17828 The following example returns the current scheduling priority for the process ID returned by the  
 17829 call to *getpid()*.

```
17830 #include <sys/resource.h>
17831 ...
17832 int which = PRIO_PROCESS;
17833 id_t pid;
17834 int ret;

17835 pid = getpid();
17836 ret = getpriority(which, pid);
```

##### 17837 Using *setpriority()*

17838 The following example sets the priority for the current process ID to  $-20$ .

```
17839 #include <sys/resource.h>
17840 ...
17841 int which = PRIO_PROCESS;
17842 id_t pid;
17843 int priority = -20;
17844 int ret;

17845 pid = getpid();
17846 ret = setpriority(which, pid, priority);
```

#### 17847 APPLICATION USAGE

17848 The *getpriority()* and *setpriority()* functions work with an offset nice value (nice value  
 17849  $-\{\text{NZERO}\}$ ). The nice value is in the range  $[0, 2*\{\text{NZERO}\} - 1]$ , while the return value for  
 17850 *getpriority()* and the third parameter for *setpriority()* are in the range  $[-\{\text{NZERO}\}, \{\text{NZERO}\} - 1]$ .

#### 17851 RATIONALE

17852 None.

#### 17853 FUTURE DIRECTIONS

17854 None.

#### 17855 SEE ALSO

17856 *nice()*, *sched\_get\_priority\_max()*, *sched\_setscheduler()*, the Base Definitions volume of  
 17857 IEEE Std 1003.1-200x, <sys/resource.h>

17858 **CHANGE HISTORY**

17859           First released in Issue 4, Version 2.

17860 **Issue 5**

17861           Moved from X/OPEN UNIX extension to BASE.

17862           The DESCRIPTION is reworded in terms of the nice value rather than *priority* to avoid confusion  
17863           with functionality in the POSIX Realtime Extension.

17864 **NAME**

17865       getprotobyname — network protocol database functions

17866 **SYNOPSIS**

17867       #include <netdb.h>

17868       struct protoent \*getprotobyname(const char \*name);

17869 **DESCRIPTION**

17870       Refer to *endprotoent()*.

17871 **NAME**

17872       getprotobynumber — network protocol database functions

17873 **SYNOPSIS**

17874       #include <netdb.h>

17875       struct protoent \*getprotobynumber(int *proto*);

17876 **DESCRIPTION**

17877       Refer to *endprotoent()*.

17878 **NAME**

17879           getprotoent — network protocol database functions

17880 **SYNOPSIS**

17881           #include <netdb.h>

17882           struct protoent \*getprotoent(void);

17883 **DESCRIPTION**

17884           Refer to *endprotoent()*.

17885 **NAME**

17886           getpwent — get user database entry

17887 **SYNOPSIS**

17888 xSI       #include <pwd.h>

17889           struct passwd \*getpwent(void);

17890

17891 **DESCRIPTION**

17892           Refer to *endpwent()*.

## 17893 NAME

17894 getpwnam, getpwnam\_r — search user database for a name

## 17895 SYNOPSIS

17896 #include <pwd.h>

17897 struct passwd \*getpwnam(const char \*name);

17898 TSF int getpwnam\_r(const char \*name, struct passwd \*pwd, char \*buffer,  
17899 size\_t bufsize, struct passwd \*\*result);

17900

## 17901 DESCRIPTION

17902 The *getpwnam()* function shall search the user database for an entry with a matching *name*.

17903 The *getpwnam()* function need not be reentrant. A function that is not required to be reentrant is  
17904 not required to be thread-safe.

17905 Applications wishing to check for error situations should set *errno* to 0 before calling  
17906 *getpwnam()*. If *getpwnam()* returns a null pointer and *errno* is non-zero, an error occurred.

17907 TSF The *getpwnam\_r()* function shall update the **passwd** structure pointed to by *pwd* and store a  
17908 pointer to that structure at the location pointed to by *result*. The structure shall contain an entry  
17909 from the user database with a matching *name*. Storage referenced by the structure is allocated  
17910 from the memory provided with the *buffer* parameter, which is *bufsize* bytes in size. The  
17911 maximum size needed for this buffer can be determined with the `{_SC_GETPW_R_SIZE_MAX}`  
17912 *sysconf()* parameter. A NULL pointer shall be returned at the location pointed to by *result* on  
17913 error or if the requested entry is not found.

## 17914 RETURN VALUE

17915 The *getpwnam()* function shall return a pointer to a **struct passwd** with the structure as defined  
17916 in <pwd.h> with a matching entry if found. A null pointer shall be returned if the requested  
17917 entry is not found, or an error occurs. On error, *errno* shall be set to indicate the error.

17918 The return value may point to a static area which is overwritten by a subsequent call to  
17919 *getpwent()*, *getpwnam()*, or *getpwuid()*.

17920 TSF If successful, the *getpwnam\_r()* function shall return zero; otherwise, an error number shall be  
17921 returned to indicate the error.

## 17922 ERRORS

17923 The *getpwnam()* and *getpwnam\_r()* functions may fail if:

17924 [EIO] An I/O error has occurred.

17925 [EINTR] A signal was caught during *getpwnam()*.

17926 [EMFILE] {OPEN\_MAX} file descriptors are currently open in the calling process.

17927 [ENFILE] The maximum allowable number of files is currently open in the system.

17928 The *getpwnam\_r()* function may fail if:

17929 TSF [ERANGE] Insufficient storage was supplied via *buffer* and *bufsize* to contain the data to  
17930 be referenced by the resulting **passwd** structure.

17931 **EXAMPLES**17932 **Getting an Entry for the Login Name**

17933 The following example uses the *getlogin()* function to return the name of the user who logged in;  
17934 this information is passed to the *getpwnam()* function to get the user database entry for that user.

```
17935 #include <sys/types.h>
17936 #include <pwd.h>
17937 #include <unistd.h>
17938 #include <stdio.h>
17939 #include <stdlib.h>
17940 ...
17941 char *lgn;
17942 struct passwd *pw;
17943 ...
17944 if ((lgn = getlogin()) == NULL || (pw = getpwnam(lgn)) == NULL) {
17945     fprintf(stderr, "Get of user information failed.\n"); exit(1);
17946 }
17947 ...
```

17948 **APPLICATION USAGE**

17949 Three names associated with the current process can be determined: *getpwuid(geteuid())* returns  
17950 the name associated with the effective user ID of the process; *getlogin()* returns the name  
17951 associated with the current login activity; and *getpwuid(getuid())* returns the name associated  
17952 with the real user ID of the process.

17953 The *getpwnam\_r()* function is thread-safe and returns values in a user-supplied buffer instead of  
17954 possibly using a static data area that may be overwritten by each call.

17955 **RATIONALE**

17956 None.

17957 **FUTURE DIRECTIONS**

17958 None.

17959 **SEE ALSO**

17960 *getpwuid()*, the Base Definitions volume of IEEE Std 1003.1-200x, <limits.h>, <pwd.h>,  
17961 <sys/types.h>

17962 **CHANGE HISTORY**

17963 First released in Issue 1. Derived from System V Release 2.0.

17964 **Issue 5**

17965 Normative text previously in the APPLICATION USAGE section is moved to the RETURN  
17966 VALUE section.

17967 The *getpwnam\_r()* function is included for alignment with the POSIX Threads Extension.

17968 A note indicating that the *getpwnam()* function need not be reentrant is added to the  
17969 DESCRIPTION.

17970 **Issue 6**

17971 The *getpwnam\_r()* function is marked as part of the Thread-Safe Functions option.

17972 The Open Group Corrigendum U028/3 is applied, correcting text in the DESCRIPTION  
17973 describing matching the *name*.



- 17974 In the SYNOPSIS, the optional include of the `<sys/types.h>` header is removed.
- 17975 In the DESCRIPTION, the note about reentrancy is expanded to cover thread-safety.
- 17976 The following new requirements on POSIX implementations derive from alignment with the  
17977 Single UNIX Specification:
- 17978 • The requirement to include `<sys/types.h>` has been removed. Although `<sys/types.h>` was  
17979 required for conforming implementations of previous POSIX specifications, it was not  
17980 required for UNIX applications.
  - 17981 • In the RETURN VALUE section, the requirement to set *errno* on error is added.
  - 17982 • The [EMFILE], [ENFILE], and [ENXIO] optional error conditions are added.
- 17983 The APPLICATION USAGE section is updated to include a note on the thread-safe function and  
17984 its avoidance of possibly using a static data area.
- 17985 IEEE PASC Interpretation 1003.1 #116 is applied, changing the description of the size of the  
17986 buffer from *bufsize* characters to bytes.

## 17987 NAME

17988 getpwuid, getpwuid\_r — search user database for a user ID

## 17989 SYNOPSIS

17990 #include &lt;pwd.h&gt;

17991 struct passwd \*getpwuid(uid\_t uid);

17992 TSF int getpwuid\_r(uid\_t uid, struct passwd \*pwd, char \*buffer,

17993 size\_t bufsize, struct passwd \*\*result);

17994

## 17995 DESCRIPTION

17996 The *getpwuid()* function shall search the user database for an entry with a matching *uid*.17997 The *getpwuid()* function need not be reentrant. A function that is not required to be reentrant is  
17998 not required to be thread-safe.17999 Applications wishing to check for error situations should set *errno* to 0 before calling *getpwuid()*.  
18000 If *getpwuid()* returns a null pointer and *errno* is set to non-zero, an error occurred.18001 TSF The *getpwuid\_r()* function shall update the **passwd** structure pointed to by *pwd* and store a |  
18002 pointer to that structure at the location pointed to by *result*. The structure shall contain an entry |  
18003 from the user database with a matching *uid*. Storage referenced by the structure is allocated  
18004 from the memory provided with the *buffer* parameter, which is *bufsize* bytes in size. The  
18005 maximum size needed for this buffer can be determined with the `{_SC_GETPW_R_SIZE_MAX}`  
18006 *sysconf()* parameter. A NULL pointer shall be returned at the location pointed to by *result* on  
18007 error or if the requested entry is not found.

## 18008 RETURN VALUE

18009 The *getpwuid()* function shall return a pointer to a **struct passwd** with the structure as defined in  
18010 <**pwd.h**> with a matching entry if found. A null pointer shall be returned if the requested entry  
18011 is not found, or an error occurs. On error, *errno* shall be set to indicate the error.18012 The return value may point to a static area which is overwritten by a subsequent call to  
18013 *getpwent()*, *getpwnam()*, or *getpwuid()*.18014 TSF If successful, the *getpwuid\_r()* function shall return zero; otherwise, an error number shall be  
18015 returned to indicate the error.

## 18016 ERRORS

18017 The *getpwuid()* and *getpwuid\_r()* functions may fail if:

18018 [EIO] An I/O error has occurred.

18019 [EINTR] A signal was caught during *getpwuid()*.

18020 [EMFILE] {OPEN\_MAX} file descriptors are currently open in the calling process.

18021 [ENFILE] The maximum allowable number of files is currently open in the system.

18022 The *getpwuid\_r()* function may fail if:18023 TSF [ERANGE] Insufficient storage was supplied via *buffer* and *bufsize* to contain the data to  
18024 be referenced by the resulting **passwd** structure.

18025 **EXAMPLES**18026 **Getting an Entry for the Root User**

18027 The following example gets the user database entry for the user with user ID 0 (root).

```
18028 #include <sys/types.h>
18029 #include <pwd.h>
18030 ...
18031 uid_t id = 0;
18032 struct passwd *pwd;
18033 pwd = getpwuid(id);
```

18034 **Finding the Name for the Effective User ID**

18035 The following example defines *pws* as a pointer to a structure of type **passwd**, which is used to  
 18036 store the structure pointer returned by the call to the *getpwuid()* function. The *geteuid()* function  
 18037 shall return the effective user ID of the calling process; this is used as the search criteria for the  
 18038 *getpwuid()* function. The call to *getpwuid()* shall return a pointer to the structure containing that  
 18039 user ID value.

```
18040 #include <unistd.h>
18041 #include <sys/types.h>
18042 #include <pwd.h>
18043 ...
18044 struct passwd *pws;
18045 pws = getpwuid(geteuid());
```

18046 **Finding an Entry in the User Database**

18047 The following example uses *getpwuid()* to search the user database for a user ID that was  
 18048 previously stored in a **stat** structure, then prints out the user name if it is found. If the user is not  
 18049 found, the program prints the numeric value of the user ID for the entry.

```
18050 #include <sys/types.h>
18051 #include <pwd.h>
18052 #include <stdio.h>
18053 ...
18054 struct stat statbuf;
18055 struct passwd *pwd;
18056 ...
18057 if ((pwd = getpwuid(statbuf.st_uid)) != NULL)
18058     printf(" %-8.8s", pwd->pw_name);
18059 else
18060     printf(" %-8d", statbuf.st_uid);
```

18061 **APPLICATION USAGE**

18062 Three names associated with the current process can be determined: *getpwuid(geteuid())* returns  
 18063 the name associated with the effective user ID of the process; *getlogin()* returns the name  
 18064 associated with the current login activity; and *getpwuid(getuid())* returns the name associated  
 18065 with the real user ID of the process.

18066 The *getpwuid\_r()* function is thread-safe and returns values in a user-supplied buffer instead of  
 18067 possibly using a static data area that may be overwritten by each call.

18068 **RATIONALE**

18069 None.

18070 **FUTURE DIRECTIONS**

18071 None.

18072 **SEE ALSO**

18073 *getpwnam()*, *geteuid()*, *getuid()*, *getlogin()*, the Base Definitions volume of IEEE Std 1003.1-200x,  
18074 `<limits.h>`, `<pwd.h>`, `<sys/types.h>`

18075 **CHANGE HISTORY**

18076 First released in Issue 1. Derived from System V Release 2.0.

18077 **Issue 5**

18078 Normative text previously in the APPLICATION USAGE section is moved to the RETURN  
18079 VALUE section.

18080 The *getpwuid\_r()* function is included for alignment with the POSIX Threads Extension.

18081 A note indicating that the *getpwuid()* function need not be reentrant is added to the  
18082 DESCRIPTION.

18083 **Issue 6**18084 The *getpwuid\_r()* function is marked as part of the Thread-Safe Functions option.

18085 The Open Group Corrigendum U028/3 is applied, correcting text in the DESCRIPTION  
18086 describing matching the *uid*.

18087 In the SYNOPSIS, the optional include of the `<sys/types.h>` header is removed.

18088 In the DESCRIPTION, the note about reentrancy is expanded to cover thread-safety.

18089 The following new requirements on POSIX implementations derive from alignment with the  
18090 Single UNIX Specification:

18091 • The requirement to include `<sys/types.h>` has been removed. Although `<sys/types.h>` was  
18092 required for conforming implementations of previous POSIX specifications, it was not  
18093 required for UNIX applications.

18094 • In the RETURN VALUE section, the requirement to set *errno* on error is added.

18095 • The [EIO], [EINTR], [EMFILE], and [ENFILE] optional error conditions are added.

18096 The APPLICATION USAGE section is updated to include a note on the thread-safe function and  
18097 its avoidance of possibly using a static data area.

18098 IEEE PASC Interpretation 1003.1 #116 is applied, changing the description of the size of the  
18099 buffer from *bufsize* characters to bytes.

## 18100 NAME

18101 getrlimit, setrlimit — control maximum resource consumption

## 18102 SYNOPSIS

18103 xsl `#include <sys/resource.h>`18104 `int getrlimit(int resource, struct rlimit *rlp);`18105 `int setrlimit(int resource, const struct rlimit *rlp);`

18106

## 18107 DESCRIPTION

18108 The *getrlimit()* function shall get, and the *setrlimit()* function shall set, limits on the consumption |  
18109 of a variety of resources. |18110 Each call to either *getrlimit()* or *setrlimit()* identifies a specific resource to be operated upon as |  
18111 well as a resource limit. A resource limit is represented by an **rlimit** structure. The *rlim\_cur* |  
18112 member specifies the current or soft limit and the *rlim\_max* member specifies the maximum or |  
18113 hard limit. Soft limits may be changed by a process to any value that is less than or equal to the |  
18114 hard limit. A process may (irreversibly) lower its hard limit to any value that is greater than or |  
18115 equal to the soft limit. Only a process with appropriate privileges can raise a hard limit. Both |  
18116 hard and soft limits can be changed in a single call to *setrlimit()* subject to the constraints |  
18117 described above.18118 The value RLIM\_INFINITY, defined in `<sys/resource.h>`, shall be considered to be larger than |  
18119 any other limit value. If a call to *getrlimit()* returns RLIM\_INFINITY for a resource, it means the |  
18120 implementation shall not enforce limits on that resource. Specifying RLIM\_INFINITY as any |  
18121 resource limit value on a successful call to *setrlimit()* shall inhibit enforcement of that resource |  
18122 limit.

18123 The following resources are defined:

18124 RLIMIT\_CORE This is the maximum size of a core file, in bytes, that may be created by a |  
18125 process. A limit of 0 shall prevent the creation of a core file. If this limit is |  
18126 exceeded, the writing of a core file shall terminate at this size.18127 RLIMIT\_CPU This is the maximum amount of CPU time, in seconds, used by a process. |  
18128 If this limit is exceeded, SIGXCPU shall be generated for the process. If |  
18129 the process is catching or ignoring SIGXCPU, or all threads belonging to |  
18130 that process are blocking SIGXCPU, the behavior is unspecified.18131 RLIMIT\_DATA This is the maximum size of a process' data segment, in bytes. If this limit |  
18132 is exceeded, the *malloc()* function shall fail with *errno* set to [ENOMEM].18133 RLIMIT\_FSIZE This is the maximum size of a file, in bytes, that may be created by a |  
18134 process. If a write or truncate operation would cause this limit to be |  
18135 exceeded, SIGXFSZ shall be generated for the thread. If the thread is |  
18136 blocking, or the process is catching or ignoring SIGXFSZ, continued |  
18137 attempts to increase the size of a file from end-of-file to beyond the limit |  
18138 shall fail with *errno* set to [EFBIG].18139 RLIMIT\_NOFILE This is a number one greater than the maximum value that the system |  
18140 may assign to a newly-created descriptor. If this limit is exceeded, |  
18141 functions that allocate new file descriptors may fail with *errno* set to |  
18142 [EMFILE]. This limit constrains the number of file descriptors that a |  
18143 process may allocate.18144 RLIMIT\_STACK This is the maximum size of a process' stack, in bytes. The |  
18145 implementation does not automatically grow the stack beyond this limit.

18146 If this limit is exceeded, SIGSEGV shall be generated for the thread. If the  
 18147 thread is blocking SIGSEGV, or the process is ignoring or catching  
 18148 SIGSEGV and has not made arrangements to use an alternate stack, the  
 18149 disposition of SIGSEGV shall be set to SIG\_DFL before it is generated.

18150 **RLIMIT\_AS** This is the maximum size of a process' total available memory, in bytes. If  
 18151 this limit is exceeded, the *malloc()* and *mmap()* functions shall fail with  
 18152 *errno* set to [ENOMEM]. In addition, the automatic stack growth fails  
 18153 with the effects outlined above.

18154 When using the *getrlimit()* function, if a resource limit can be represented correctly in an object  
 18155 of type **rlim\_t**, then its representation is returned; otherwise, if the value of the resource limit is  
 18156 equal to that of the corresponding saved hard limit, the value returned shall be  
 18157 RLIM\_SAVED\_MAX; otherwise, the value returned shall be RLIM\_SAVED\_CUR.

18158 When using the *setrlimit()* function, if the requested new limit is RLIM\_INFINITY, the new limit  
 18159 shall be “no limit”; otherwise, if the requested new limit is RLIM\_SAVED\_MAX, the new limit  
 18160 shall be the corresponding saved hard limit; otherwise, if the requested new limit is  
 18161 RLIM\_SAVED\_CUR, the new limit shall be the corresponding saved soft limit; otherwise, the  
 18162 new limit shall be the requested value. In addition, if the corresponding saved limit can be  
 18163 represented correctly in an object of type **rlim\_t** then it shall be overwritten with the new limit.

18164 The result of setting a limit to RLIM\_SAVED\_MAX or RLIM\_SAVED\_CUR is unspecified unless  
 18165 a previous call to *getrlimit()* returned that value as the soft or hard limit for the corresponding  
 18166 resource limit.

18167 The determination of whether a limit can be correctly represented in an object of type **rlim\_t** is  
 18168 implementation-defined. For example, some implementations permit a limit whose value is  
 18169 greater than RLIM\_INFINITY and others do not.

18170 The *exec* family of functions shall cause resource limits to be saved. |

18171 **RETURN VALUE**

18172 Upon successful completion, *getrlimit()* and *setrlimit()* shall return 0. Otherwise, these functions  
 18173 shall return -1 and set *errno* to indicate the error.

18174 **ERRORS**

18175 The *getrlimit()* and *setrlimit()* functions shall fail if:

18176 [EINVAL] An invalid *resource* was specified; or in a *setrlimit()* call, the new *rlim\_cur*  
 18177 exceeds the new *rlim\_max*.

18178 [EPERM] The limit specified to *setrlimit()* would have raised the maximum limit value,  
 18179 and the calling process does not have appropriate privileges.

18180 The *setrlimit()* function may fail if:

18181 [EINVAL] The limit specified cannot be lowered because current usage is already higher  
 18182 than the limit.

18183 **EXAMPLES**

18184           None.

18185 **APPLICATION USAGE**

18186           If a process attempts to set the hard limit or soft limit for RLIMIT\_NOFILE to less than the value  
18187           of `{_POSIX_OPEN_MAX}` from `<limits.h>`, unexpected behavior may occur.

18188 **RATIONALE**

18189           None.

18190 **FUTURE DIRECTIONS**

18191           None.

18192 **SEE ALSO**

18193           *exec*, *fork()*, *malloc()*, *open()*, *sigaltstack()*, *sysconf()*, *ulimit()*, the Base Definitions volume of  
18194           IEEE Std 1003.1-200x, `<stropts.h>`, `<sys/resource.h>`

18195 **CHANGE HISTORY**

18196           First released in Issue 4, Version 2.

18197 **Issue 5**

18198           Moved from X/OPEN UNIX extension to BASE and an APPLICATION USAGE section is added.

18199           Large File Summit extensions are added.

18200 **NAME**

18201 getrusage — get information about resource utilization

18202 **SYNOPSIS**18203 XSI `#include <sys/resource.h>`18204 `int getrusage(int who, struct rusage *r_usage);`

18205

18206 **DESCRIPTION**

18207 The `getrusage()` function shall provide measures of the resources used by the current process or  
 18208 its terminated and waited-for child processes. If the value of the `who` argument is  
 18209 `RUSAGE_SELF`, information shall be returned about resources used by the current process. If the  
 18210 value of the `who` argument is `RUSAGE_CHILDREN`, information shall be returned about  
 18211 resources used by the terminated and waited-for children of the current process. If the child is  
 18212 never waited for (for example, if the parent has `SA_NOCLDWAIT` set or sets `SIGCHLD` to  
 18213 `SIG_IGN`), the resource information for the child process is discarded and not included in the  
 18214 resource information provided by `getrusage()`.

18215 The `r_usage` argument is a pointer to an object of type **struct rusage** in which the returned  
 18216 information is stored.

18217 **RETURN VALUE**

18218 Upon successful completion, `getrusage()` shall return 0; otherwise, `-1` shall be returned and `errno`  
 18219 set to indicate the error.

18220 **ERRORS**18221 The `getrusage()` function shall fail if:18222 `[EINVAL]` The value of the `who` argument is not valid.18223 **EXAMPLES**18224 **Using getrusage()**

18225 The following example returns information about the resources used by the current process.

18226 `#include <sys/resource.h>`18227 `...`18228 `int who = RUSAGE_SELF;`18229 `struct rusage usage;`18230 `int ret;`18231 `ret = getrusage(who, &usage);`18232 **APPLICATION USAGE**

18233 None.

18234 **RATIONALE**

18235 None.

18236 **FUTURE DIRECTIONS**

18237 None.

18238 **SEE ALSO**18239 `exit()`, `sigaction()`, `time()`, `times()`, `wait()`, the Base Definitions volume of IEEE Std 1003.1-200x,18240 `<sys/resource.h>`



18241 **CHANGE HISTORY**

18242 First released in Issue 4, Version 2.

18243 **Issue 5**

18244 Moved from X/OPEN UNIX extension to BASE.

18245 **NAME**

18246 gets — get a string from a stdin stream

18247 **SYNOPSIS**

18248 #include &lt;stdio.h&gt;

18249 char \*gets(char \*s);

18250 **DESCRIPTION**

18251 cx The functionality described on this reference page is aligned with the ISO C standard. Any  
18252 conflict between the requirements described here and the ISO C standard is unintentional. This  
18253 volume of IEEE Std 1003.1-200x defers to the ISO C standard.

18254 The *gets()* function shall read bytes from the standard input stream, *stdin*, into the array pointed  
18255 to by *s*, until a newline is read or an end-of-file condition is encountered. Any <newline> shall  
18256 be discarded and a null byte shall be placed immediately after the last byte read into the array.

18257 cx The *gets()* function may mark the *st\_atime* field of the file associated with *stream* for update. The  
18258 *st\_atime* field shall be marked for update by the first successful execution of *fgetc()*, *fgets()*,  
18259 *fread()*, *getc()*, *getchar()*, *gets()*, *fscanf()*, or *scanf()* using *stream* that returns data not supplied by  
18260 a prior call to *ungetc()*.

18261 **RETURN VALUE**

18262 Upon successful completion, *gets()* shall return *s*. If the stream is at end-of-file, the end-of-file  
18263 indicator for the stream shall be set and *gets()* shall return a null pointer. If a read error occurs,  
18264 cx the error indicator for the stream shall be set, *gets()* shall return a null pointer and set *errno* to  
18265 indicate the error.

18266 **ERRORS**18267 Refer to *fgetc()*.18268 **EXAMPLES**

18269 None.

18270 **APPLICATION USAGE**

18271 Reading a line that overflows the array pointed to by *s* results in undefined behavior. The use of  
18272 *fgets()* is recommended.

18273 Since the user cannot specify the length of the buffer passed to *gets()*, use of this function is  
18274 discouraged. The length of the string read is unlimited. It is possible to overflow this buffer in  
18275 such a way as to cause applications to fail, or possible system security violations.

18276 It is recommended that the *fgets()* function should be used to read input lines.

18277 **RATIONALE**

18278 None.

18279 **FUTURE DIRECTIONS**

18280 None.

18281 **SEE ALSO**18282 *feof()*, *ferror()*, *fgets()*, the Base Definitions volume of IEEE Std 1003.1-200x, <stdio.h>18283 **CHANGE HISTORY**

18284 First released in Issue 1. Derived from Issue 1 of the SVID.

18285 **Issue 6**

18286 Extensions beyond the ISO C standard are now marked.

18287 **NAME**

18288        getservbyname — network services database functions

18289 **SYNOPSIS**

18290        #include <netdb.h>

18291        struct servent \*getservbyname(const char \*name, const char \*proto);

18292 **DESCRIPTION**

18293        Refer to *endservent()*.

18294 **NAME**

18295        getservbyport — network services database functions

18296 **SYNOPSIS**

18297        #include <netdb.h>

18298        struct servent \*getservbyport(int *port*, const char \**proto*);

18299 **DESCRIPTION**

18300        Refer to *endservent()*.

18301 **NAME**

18302        getservent — network services database functions

18303 **SYNOPSIS**

18304        #include <netdb.h>

18305        struct servent \*getservent(void);

18306 **DESCRIPTION**

18307        Refer to *endservent()*.

18308 **NAME**

18309        getsid — get the process group ID of a session leader

18310 **SYNOPSIS**

18311 XSI        #include <unistd.h>

18312        pid\_t getsid(pid\_t pid);

18313

18314 **DESCRIPTION**

18315        The *getsid()* function shall obtain the process group ID of the process that is the session leader of  
18316        the process specified by *pid*. If *pid* is (**pid\_t**)0, it specifies the calling process.

18317 **RETURN VALUE**

18318        Upon successful completion, *getsid()* shall return the process group ID of the session leader of  
18319        the specified process. Otherwise, it shall return (**pid\_t**)-1 and set *errno* to indicate the error.

18320 **ERRORS**

18321        The *getsid()* function shall fail if:

18322        [EPERM]        The process specified by *pid* is not in the same session as the calling process,  
18323        and the implementation does not allow access to the process group ID of the  
18324        session leader of that process from the calling process.

18325        [ESRCH]        There is no process with a process ID equal to *pid*.

18326 **EXAMPLES**

18327        None.

18328 **APPLICATION USAGE**

18329        None.

18330 **RATIONALE**

18331        None.

18332 **FUTURE DIRECTIONS**

18333        None.

18334 **SEE ALSO**

18335        *exec*, *fork()*, *getpid()*, *getpgid()*, *setpgid()*, *setsid()*, the Base Definitions volume of  
18336        IEEE Std 1003.1-200x, <unistd.h>

18337 **CHANGE HISTORY**

18338        First released in Issue 4, Version 2.

18339 **Issue 5**

18340        Moved from X/OPEN UNIX extension to BASE.

18341 **NAME**

18342 getsockname — get the socket name

18343 **SYNOPSIS**

18344 #include &lt;sys/socket.h&gt;

18345 int getsockname(int *socket*, struct sockaddr \*restrict *address*,  
18346 socklen\_t \*restrict *address\_len*);18347 **DESCRIPTION**18348 The *getsockname()* function shall retrieve the locally-bound name of the specified socket, store  
18349 this address in the **sockaddr** structure pointed to by the *address* argument, and store the length of  
18350 this address in the object pointed to by the *address\_len* argument.18351 If the actual length of the address is greater than the length of the supplied **sockaddr** structure,  
18352 the stored address shall be truncated.18353 If the socket has not been bound to a local name, the value stored in the object pointed to by  
18354 *address* is unspecified.18355 **RETURN VALUE**18356 Upon successful completion, 0 shall be returned, the *address* argument shall point to the address  
18357 of the socket, and the *address\_len* argument shall point to the length of the address. Otherwise, -1  
18358 shall be returned and *errno* set to indicate the error.18359 **ERRORS**18360 The *getsockname()* function shall fail if:18361 [EBADF] The *socket* argument is not a valid file descriptor.18362 [ENOTSOCK] The *socket* argument does not refer to a socket.

18363 [EOPNOTSUPP] The operation is not supported for this socket's protocol.

18364 The *getsockname()* function may fail if:

18365 [EINVAL] The socket has been shut down.

18366 [ENOBUFS] Insufficient resources were available in the system to complete the function.

18367 **EXAMPLES**

18368 None.

18369 **APPLICATION USAGE**

18370 None.

18371 **RATIONALE**

18372 None.

18373 **FUTURE DIRECTIONS**

18374 None.

18375 **SEE ALSO**18376 *accept()*, *bind()*, *getpeername()*, *socket()*, the Base Definitions volume of IEEE Std 1003.1-200x,  
18377 <sys/socket.h>18378 **CHANGE HISTORY**

18379 First released in Issue 6. Derived from the XNS, Issue 5.2 specification.

18380 The **restrict** keyword is added to the *getsockname()* prototype for alignment with the  
18381 ISO/IEC 9899:1999 standard.

## 18382 NAME

18383 getsockopt — get the socket options

## 18384 SYNOPSIS

18385 #include &lt;sys/socket.h&gt;

```
18386 int getsockopt(int socket, int level, int option_name,
18387               void *restrict option_value, socklen_t *restrict option_len);
```

## 18388 DESCRIPTION

18389 The *getsockopt()* function manipulates options associated with a socket.

18390 The *getsockopt()* function shall retrieve the value for the option specified by the *option\_name*  
 18391 argument for the socket specified by the *socket* argument. If the size of the option value is greater  
 18392 than *option\_len*, the value stored in the object pointed to by the *option\_value* argument shall be  
 18393 silently truncated. Otherwise, the object pointed to by the *option\_len* argument shall be modified  
 18394 to indicate the actual length of the value.

18395 The *level* argument specifies the protocol level at which the option resides. To retrieve options at  
 18396 the socket level, specify the *level* argument as SOL\_SOCKET. To retrieve options at other levels,  
 18397 supply the appropriate level identifier for the protocol controlling the option. For example, to  
 18398 indicate that an option is interpreted by the TCP (Transmission Control Protocol), set *level* to  
 18399 IPPROTO\_TCP as defined in the <netinet/in.h> header.

18400 The socket in use may require the process to have appropriate privileges to use the *getsockopt()*  
 18401 function.

18402 The *option\_name* argument specifies a single option to be retrieved. It can be one of the following  
 18403 values defined in <sys/socket.h>:

18404	SO_DEBUG	Reports whether debugging information is being recorded. This option	
18405		shall store an <b>int</b> value. This is a Boolean option.	
18406	SO_ACCEPTCONN	Reports whether socket listening is enabled. This option shall store an <b>int</b>	
18407		value. This is a Boolean option.	
18408	SO_BROADCAST	Reports whether transmission of broadcast messages is supported, if this	
18409		is supported by the protocol. This option shall store an <b>int</b> value. This is a	
18410		Boolean option.	
18411	SO_REUSEADDR	Reports whether the rules used in validating addresses supplied to <i>bind()</i>	
18412		should allow reuse of local addresses, if this is supported by the protocol.	
18413		This option shall store an <b>int</b> value. This is a Boolean option.	
18414	SO_KEEPALIVE	Reports whether connections are kept active with periodic transmission	
18415		of messages, if this is supported by the protocol.	
18416		If the connected socket fails to respond to these messages, the connection	
18417		shall be broken and threads writing to that socket shall be notified with a	
18418		SIGPIPE signal. This option shall store an <b>int</b> value. This is a Boolean	
18419		option.	
18420	SO_LINGER	Reports whether the socket lingers on <i>close()</i> if data is present. If	
18421		SO_LINGER is set, the system blocks the process during <i>close()</i> until it	
18422		can transmit the data or until the end of the interval indicated by the	
18423		<i>l_linger</i> member, whichever comes first. If SO_LINGER is not specified,	
18424		and <i>close()</i> is issued, the system handles the call in a way that allows the	
18425		process to continue as quickly as possible. This option shall store a <b>linger</b>	
18426		structure.	



18427	SO_OOBINLINE	Reports whether the socket leaves received out-of-band data (data marked urgent) inline. This option shall store an <b>int</b> value. This is a Boolean option.
18428		
18429		
18430	SO_SNDBUF	Reports send buffer size information. This option shall store an <b>int</b> value.
18431	SO_RCVBUF	Reports receive buffer size information. This option shall store an <b>int</b> value.
18432		
18433	SO_ERROR	Reports information about error status and clears it. This option shall store an <b>int</b> value.
18434		
18435	SO_TYPE	Reports the socket type. This option shall store an <b>int</b> value. Socket types are described in Section 2.10.6 (on page 509).
18436		
18437	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 shall store an <b>int</b> value. This is a Boolean option.
18438		
18439		
18440		
18441		
18442		
18443	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 from that returned; for example, out-of-band data.) This option shall store an <b>int</b> value. Note that not all implementations allow this option to be retrieved.
18444		
18445		
18446		
18447		
18448		
18449		
18450		
18451		
18452	SO_RCVTIMEO	Reports the timeout value for input operations. This option shall store 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.
18453		
18454		
18455		
18456		
18457		
18458		
18459		
18460	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 shall store an <b>int</b> value. Note that not all implementations allow this option to be retrieved.
18461		
18462		
18463		
18464		
18465	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 shall store a <b>timeval</b> structure. Note that not all implementations allow this option to be retrieved.
18466		
18467		
18468		
18469		
18470		
18471		
18472		
18473		
		For Boolean options, a zero value indicates that the option is disabled and a non-zero value indicates that the option is enabled.

- 18474 Options at other protocol levels vary in format and name.
- 18475 The socket in use may require the process to have appropriate privileges to use the *getsockopt()*  
18476 function.
- 18477 **RETURN VALUE**
- 18478 Upon successful completion, *getsockopt()* shall return 0; otherwise, -1 shall be returned and *errno*  
18479 set to indicate the error.
- 18480 **ERRORS**
- 18481 The *getsockopt()* function shall fail if:
- 18482 [EBADF] The *socket* argument is not a valid file descriptor.
- 18483 [EINVAL] The specified option is invalid at the specified socket level.
- 18484 [ENOPROTOOPT]
- 18485 The option is not supported by the protocol.
- 18486 [ENOTSOCK] The *socket* argument does not refer to a socket.
- 18487 The *getsockopt()* function may fail if:
- 18488 [EACCES] The calling process does not have the appropriate privileges.
- 18489 [EINVAL] The socket has been shut down.
- 18490 [ENOBUFS] Insufficient resources are available in the system to complete the function.
- 18491 **EXAMPLES**
- 18492 None.
- 18493 **APPLICATION USAGE**
- 18494 None.
- 18495 **RATIONALE**
- 18496 None.
- 18497 **FUTURE DIRECTIONS**
- 18498 None.
- 18499 **SEE ALSO**
- 18500 *bind()*, *close()*, *endprotoent()*, *setsockopt()*, *socket()*, the Base Definitions volume of  
18501 IEEE Std 1003.1-200x, <sys/socket.h>, <netinet/in.h>
- 18502 **CHANGE HISTORY**
- 18503 First released in Issue 6. Derived from the XNS, Issue 5.2 specification.
- 18504 The **restrict** keyword is added to the *getsockopt()* prototype for alignment with the  
18505 ISO/IEC 9899:1999 standard.

18506 **NAME**

18507       getsubopt — parse suboption arguments from a string

18508 **SYNOPSIS**

18509 XSI       #include &lt;stdlib.h&gt;

18510       int getsubopt(char \*\**optionp*, char \* const \**tokens*, char \*\**valuep*);

18511

18512 **DESCRIPTION**18513       The *getsubopt()* function shall parse suboption arguments in a flag argument. Such options often  
18514       result from the use of *getopt()*.18515       The *getsubopt()* argument *optionp* is a pointer to a pointer to the option argument string. The  
18516       suboption arguments shall be separated by commas and each may consist of either a single  
18517       token, or a token-value pair separated by an equal sign.18518       The *keylistp* argument shall be a pointer to a vector of strings. The end of the vector is identified  
18519       by a null pointer. Each entry in the vector is one of the possible tokens that might be found in  
18520       \**optionp*. Since commas delimit suboption arguments in *optionp*, they should not appear in any of  
18521       the strings pointed to by *keylistp*. Similarly, because an equal sign separates a token from its  
18522       value, the application should not include an equal sign in any of the strings pointed to by  
18523       *keylistp*.18524       The *valuep* argument is the address of a value string pointer.18525       If a comma appears in *optionp*, it shall be interpreted as a suboption separator. After commas  
18526       have been processed, if there are one or more equal signs in a suboption string, the first equal  
18527       sign in any suboption string shall be interpreted as a separator between a token and a value.  
18528       Subsequent equal signs in a suboption string shall be interpreted as part of the value.18529       If the string at \**optionp* contains only one suboption argument (equivalently, no commas),  
18530       *getsubopt()* shall update \**optionp* to point to the nul character at the end of the string. Otherwise,  
18531       it shall isolate the suboption argument by replacing the comma separator with a nul character,  
18532       and shall update \**optionp* to point to the start of the next suboption argument. If the suboption  
18533       argument has an associated value (equivalently, contains an equal sign), *getsubopt()* shall update  
18534       \**valuep* to point to the value's first character. Otherwise, it shall set \**valuep* to a null pointer. The  
18535       calling application may use this information to determine whether the presence or absence of a  
18536       value for the suboption is an error.18537       Additionally, when *getsubopt()* fails to match the suboption argument with a token in the *keylistp*  
18538       array, the calling application should decide if this is an error, or if the unrecognized option  
18539       should be processed in another way.18540 **RETURN VALUE**18541       The *getsubopt()* function shall return the index of the matched token string, or -1 if no token  
18542       strings were matched.18543 **ERRORS**

18544       No errors are defined.

## 18545 EXAMPLES

```

18546     #include <stdio.h>
18547     #include <stdlib.h>

18548     int do_all;
18549     const char *type;
18550     int read_size;
18551     int write_size;
18552     int read_only;

18553     enum
18554     {
18555         RO_OPTION = 0,
18556         RW_OPTION,
18557         READ_SIZE_OPTION,
18558         WRITE_SIZE_OPTION
18559     };

18560     const char *mount_opts[] =
18561     {
18562         [RO_OPTION] = "ro",
18563         [RW_OPTION] = "rw",
18564         [READ_SIZE_OPTION] = "rsize",
18565         [WRITE_SIZE_OPTION] = "wsize",
18566         NULL
18567     };

18568     int
18569     main(int argc, char *argv[])
18570     {
18571         char *subopts, *value;
18572         int opt;

18573         while ((opt = getopt(argc, argv, "at:o:")) != -1)
18574             switch(opt)
18575             {
18576                 case 'a':
18577                     do_all = 1;
18578                     break;
18579                 case 't':
18580                     type = optarg;
18581                     break;
18582                 case 'o':
18583                     subopts = optarg;
18584                     while (*subopts != '\0')
18585                         switch(getsubopt(&subopts, mount_opts, &value))
18586                         {
18587                             case RO_OPTION:
18588                                 read_only = 1;
18589                                 break;
18590                             case RW_OPTION:
18591                                 read_only = 0;
18592                                 break;
18593                             case READ_SIZE_OPTION:

```

```

18594         if (value == NULL)
18595             abort();
18596         read_size = atoi(value);
18597         break;
18598     case WRITE_SIZE_OPTION:
18599         if (value == NULL)
18600             abort();
18601         write_size = atoi(value);
18602         break;
18603     default:
18604         /* Unknown suboption. */
18605         printf("Unknown suboption '%s'\n", value);
18606         break;
18607     }
18608     break;
18609     default:
18610         abort();
18611     }
18612     /* Do the real work. */
18613     return 0;
18614 }

```

### 18615 Parsing Suboptions

18616 The following example uses the *getsubopt()* function to parse a value argument in the *optarg*  
 18617 external variable returned by a call to *getopt()*.

```

18618 #include <stdlib.h>
18619 ...
18620 char *tokens[] = {"HOME", "PATH", "LOGNAME", (char *) NULL };
18621 char *value;
18622 int opt, index;
18623 while ((opt = getopt(argc, argv, "e:")) != -1) {
18624     switch(opt) {
18625     case 'e' :
18626         while ((index = getsubopt(&optarg, tokens, &value)) != -1) {
18627             switch(index) {
18628             ...
18629             }
18630             break;
18631         ...
18632     }
18633     }
18634     ...

```

### 18635 APPLICATION USAGE

18636 None.

### 18637 RATIONALE

18638 None.

18639 **FUTURE DIRECTIONS**

18640           None.

18641 **SEE ALSO**

18642           *getopt()*, the Base Definitions volume of IEEE Std 1003.1-200x, <**stdlib.h**>

18643 **CHANGE HISTORY**

18644           First released in Issue 4, Version 2.

18645 **Issue 5**

18646           Moved from X/OPEN UNIX extension to BASE.

18647 **NAME**

18648           gettimeofday — get the date and time

18649 **SYNOPSIS**

18650 XSI       #include &lt;sys/time.h&gt;

18651           int gettimeofday(struct timeval \*restrict tp, void \*restrict tzp);

18652

18653 **DESCRIPTION**

18654       The *gettimeofday()* function shall obtain the current time, expressed as seconds and |  
18655       microseconds since the Epoch, and store it in the **timeval** structure pointed to by *tp*. The |  
18656       resolution of the system clock is unspecified.

18657       If *tzp* is not a null pointer, the behavior is unspecified.

18658 **RETURN VALUE**

18659       The *gettimeofday()* function shall return 0 and no value shall be reserved to indicate an error.

18660 **ERRORS**

18661       No errors are defined.

18662 **EXAMPLES**

18663       None.

18664 **APPLICATION USAGE**

18665       None.

18666 **RATIONALE**

18667       None.

18668 **FUTURE DIRECTIONS**

18669       None.

18670 **SEE ALSO**

18671       *ctime()*, *ftime()*, the Base Definitions volume of IEEE Std 1003.1-200x, <sys/time.h>

18672 **CHANGE HISTORY**

18673       First released in Issue 4, Version 2.

18674 **Issue 5**

18675       Moved from X/OPEN UNIX extension to BASE.

18676 **Issue 6**

18677       The DESCRIPTION is updated to refer to “seconds since the Epoch” rather than “seconds since  
18678       00:00:00 UTC (Coordinated Universal Time), January 1 1970” for consistency with other *time*  
18679       functions.

18680       The **restrict** keyword is added to the *gettimeofday()* prototype for alignment with the  
18681       ISO/IEC 9899:1999 standard.

18682 **NAME**

18683           getuid — get a real user ID

18684 **SYNOPSIS**

18685           #include <unistd.h>

18686           uid\_t getuid(void);

18687 **DESCRIPTION**

18688           The *getuid()* function shall return the real user ID of the calling process.

18689 **RETURN VALUE**

18690           The *getuid()* function shall always be successful and no return value is reserved to indicate the  
18691           error.

18692 **ERRORS**

18693           No errors are defined.

18694 **EXAMPLES**18695           **Setting the Effective User ID to the Real User ID**

18696           The following example sets the effective user ID and the real user ID of the current process to the  
18697           real user ID of the caller.

```
18698           #include <unistd.h>
18699           #include <sys/types.h>
18700           ...
18701           setreuid(getuid(), getuid());
18702           ...
```

18703 **APPLICATION USAGE**

18704           None.

18705 **RATIONALE**

18706           None.

18707 **FUTURE DIRECTIONS**

18708           None.

18709 **SEE ALSO**

18710           *getegid()*, *geteuid()*, *getgid()*, *setegid()*, *seteuid()*, *setgid()*, *setregid()*, *setreuid()*, *setuid()*, the Base  
18711           Definitions volume of IEEE Std 1003.1-200x, <sys/types.h>, <unistd.h>

18712 **CHANGE HISTORY**

18713           First released in Issue 1. Derived from Issue 1 of the SVID.

18714 **Issue 6**

18715           In the SYNOPSIS, the optional include of the <sys/types.h> header is removed.

18716           The following new requirements on POSIX implementations derive from alignment with the  
18717           Single UNIX Specification:

- 18718           • The requirement to include <sys/types.h> has been removed. Although <sys/types.h> was  
18719           required for conforming implementations of previous POSIX specifications, it was not  
18720           required for UNIX applications.



18721 **NAME**

18722           getutxent, getutxid, getutxline — get user accounting database entries

18723 **SYNOPSIS**

18724 XSI       #include &lt;utmpx.h&gt;

18725           struct utmpx \*getutxent(void);

18726           struct utmpx \*getutxid(const struct utmpx \*id);

18727           struct utmpx \*getutxline(const struct utmpx \*line);

18728

18729 **DESCRIPTION**18730           Refer to *endutxent()*.

18731 **NAME**

18732           getwc — get a wide character from a stream

18733 **SYNOPSIS**

18734           #include <stdio.h>

18735           #include <wchar.h>

18736           wint\_t getwc(FILE \**stream*);

18737 **DESCRIPTION**

18738 **CX**       The functionality described on this reference page is aligned with the ISO C standard. Any  
18739 conflict between the requirements described here and the ISO C standard is unintentional. This  
18740 volume of IEEE Std 1003.1-200x defers to the ISO C standard.

18741       The *getwc()* function shall be equivalent to *fgetwc()*, except that if it is implemented as a macro it  
18742 may evaluate *stream* more than once, so the argument should never be an expression with side  
18743 effects.

18744 **RETURN VALUE**

18745       Refer to *fgetwc()*.

18746 **ERRORS**

18747       Refer to *fgetwc()*.

18748 **EXAMPLES**

18749       None.

18750 **APPLICATION USAGE**

18751       Since it may be implemented as a macro, *getwc()* may treat incorrectly a *stream* argument with  
18752 side effects. In particular, *getwc(\*f++)* does not necessarily work as expected. Therefore, use of  
18753 this function is not recommended; *fgetwc()* should be used instead.

18754 **RATIONALE**

18755       None.

18756 **FUTURE DIRECTIONS**

18757       None.

18758 **SEE ALSO**

18759       *fgetwc()*, the Base Definitions volume of IEEE Std 1003.1-200x, <stdio.h>, <wchar.h>

18760 **CHANGE HISTORY**

18761       First released as a World-wide Portability Interface in Issue 4. Derived from the MSE working  
18762 draft.

18763 **Issue 5**

18764       The Optional Header (OH) marking is removed from <stdio.h>.

18765 **NAME**

18766        getwchar — get a wide character from a stdin stream

18767 **SYNOPSIS**

18768        #include <wchar.h>

18769        wint\_t getwchar(void);

18770 **DESCRIPTION**

18771 cx       The functionality described on this reference page is aligned with the ISO C standard. Any  
18772       conflict between the requirements described here and the ISO C standard is unintentional. This  
18773       volume of IEEE Std 1003.1-200x defers to the ISO C standard.

18774        The *getwchar()* function shall be equivalent to *getwc(stdin)*.

18775 **RETURN VALUE**

18776        Refer to *fgetwc()*.

18777 **ERRORS**

18778        Refer to *fgetwc()*.

18779 **EXAMPLES**

18780        None.

18781 **APPLICATION USAGE**

18782        If the **wint\_t** value returned by *getwchar()* is stored into a variable of type **wchar\_t** and then  
18783        compared against the **wint\_t** macro WEOF, the result may be incorrect. Only the **wint\_t** type is  
18784        guaranteed to be able to represent any wide character and WEOF.

18785 **RATIONALE**

18786        None.

18787 **FUTURE DIRECTIONS**

18788        None.

18789 **SEE ALSO**

18790        *fgetwc()*, *getwc()*, the Base Definitions volume of IEEE Std 1003.1-200x, <**wchar.h**>

18791 **CHANGE HISTORY**

18792        First released as a World-wide Portability Interface in Issue 4. Derived from the MSE working  
18793        draft.

18794 **NAME**18795 getwd — get the current working directory pathname (**LEGACY**) |18796 **SYNOPSIS**

18797 XSI #include &lt;unistd.h&gt;

18798 char \*getwd(char \*path\_name);

18799

18800 **DESCRIPTION**

18801 The *getwd()* function shall determine an absolute pathname of the current working directory of |  
18802 the calling process, and copy a string containing that pathname into the array pointed to by the |  
18803 *path\_name* argument.

18804 If the length of the pathname of the current working directory is greater than ({PATH\_MAX}+1) |  
18805 including the null byte, *getwd()* shall fail and return a null pointer.

18806 **RETURN VALUE**

18807 Upon successful completion, a pointer to the string containing the absolute pathname of the |  
18808 current working directory shall be returned. Otherwise, *getwd()* shall return a null pointer and |  
18809 the contents of the array pointed to by *path\_name* are undefined.

18810 **ERRORS**

18811 No errors are defined.

18812 **EXAMPLES**

18813 None.

18814 **APPLICATION USAGE**

18815 For applications portability, the *getcwd()* function should be used to determine the current |  
18816 working directory instead of *getwd()*.

18817 **RATIONALE**

18818 Since the user cannot specify the length of the buffer passed to *getwd()*, use of this function is |  
18819 discouraged. The length of a pathname described in {PATH\_MAX} is file system-dependent and |  
18820 may vary from one mount point to another, or might even be unlimited. It is possible to |  
18821 overflow this buffer in such a way as to cause applications to fail, or possible system security |  
18822 violations.

18823 It is recommended that the *getcwd()* function should be used to determine the current working |  
18824 directory.

18825 **FUTURE DIRECTIONS**

18826 This function may be withdrawn in a future version.

18827 **SEE ALSO**18828 *getcwd()*, the Base Definitions volume of IEEE Std 1003.1-200x, <unistd.h>18829 **CHANGE HISTORY**

18830 First released in Issue 4, Version 2.

18831 **Issue 5**

18832 Moved from X/OPEN UNIX extension to BASE.

18833 **Issue 6**

18834 This function is marked LEGACY.

## 18835 NAME

18836 glob, globfree — generate pathnames matching a pattern |

## 18837 SYNOPSIS

18838 #include &lt;glob.h&gt;

18839 int glob(const char \*restrict *pattern*, int *flags*,18840 int(\**errfunc*)(const char \**epath*, int *eerrno*),18841 glob\_t \*restrict *pglob*);18842 void globfree(glob\_t \**pglob*);

## 18843 DESCRIPTION

18844 The *glob()* function is a pathname generator that shall implement the rules defined in the Shell |

18845 and Utilities volume of IEEE Std 1003.1-200x, Section 2.13, Pattern Matching Notation, with |

18846 optional support for rule 3 in the Shell and Utilities volume of IEEE Std 1003.1-200x, Section |

18847 2.13.3, Patterns Used for Filename Expansion. |

18848 The structure type **glob\_t** is defined in <**glob.h**> and includes at least the following members: |

18849

18850

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 pathnames.
size_t	<i>gl_offs</i>	Slots to reserve at the beginning of <i>gl_pathv</i> .

18851

18852

18853

18854 The argument *pattern* is a pointer to a pathname pattern to be expanded. The *glob()* function |

18855 shall match all accessible pathnames against this pattern and develop a list of all pathnames that |

18856 match. In order to have access to a pathname, *glob()* requires search permission on every |

18857 component of a path except the last, and read permission on each directory of any filename |

18858 component of *pattern* that contains any of the following special characters: '\*', '?', and '['.18859 The *glob()* function shall store the number of matched pathnames into *pglob->gl\_pathc* and a |18860 pointer to a list of pointers to path names into *pglob->gl\_pathv*. The pathnames shall be in sort |18861 order as defined by the current setting of the *LC\_COLLATE* category; see the Base Definitions |18862 volume of IEEE Std 1003.1-200x, Section 7.3.2, *LC\_COLLATE*. The first pointer after the last |

18863 pathname shall be a null pointer. If the pattern does not match any pathnames, the returned |

18864 number of matched paths is set to 0, and the contents of *pglob->gl\_pathv* are implementation- |

18865 defined.

18866 It is the caller's responsibility to create the structure pointed to by *pglob*. The *glob()* function shall |18867 allocate other space as needed, including the memory pointed to by *gl\_pathv*. The *globfree()* |18868 function shall free any space associated with *pglob* from a previous call to *glob()*.18869 The *flags* argument is used to control the behavior of *glob()*. The value of *flags* is a bitwise- |18870 inclusive OR of zero or more of the following constants, which are defined in <**glob.h**>: |18871 **GLOBAL\_APPEND** Append pathnames generated to the ones from a previous call to *glob()*. |18872 **GLOBAL\_DOOFFS** Make use of *pglob->gl\_offs*. If this flag is set, *pglob->gl\_offs* is used to |18873 specify how many null pointers to add to the beginning of *pglob->gl\_pathv*. In other words, |18874 *pglob->gl\_pathv* shall point to *pglob->gl\_offs* null |18875 pointers, followed by *pglob->gl\_pathc* pathname pointers, followed by a |

18876 null pointer.

18877 **GLOBAL\_ERR** Cause *glob()* to return when it encounters a directory that it cannot open |18878 or read. Ordinarily, *glob()* continues to find matches.

18879	GLOB_MARK	Each pathname that is a directory that matches <i>pattern</i> shall have a slash	
18880		appended.	
18881	GLOB_NOCHECK	Supports rule 3 in the Shell and Utilities volume of IEEE Std 1003.1-200x,	
18882		Section 2.13.3, Patterns Used for Filename Expansion. If <i>pattern</i> does not	
18883		match any pathname, then <i>glob()</i> shall return a list consisting of only	
18884		<i>pattern</i> , and the number of matched pathnames is 1.	
18885	GLOB_NOESCAPE	Disable backslash escaping.	
18886	GLOB_NOSORT	Ordinarily, <i>glob()</i> sorts the matching pathnames according to the current	
18887		setting of the <i>LC_COLLATE</i> category, see the Base Definitions volume of	
18888		IEEE Std 1003.1-200x, Section 7.3.2, <i>LC_COLLATE</i> . When this flag is	
18889		used, the order of path names returned is unspecified.	
18890	The GLOB_APPEND flag can be used to append a new set of pathnames to those found in a		
18891	previous call to <i>glob()</i> . The following rules apply to applications when two or more calls to		
18892	<i>glob()</i> are made with the same value of <i>pglob</i> and without intervening calls to <i>globfree()</i> :		
18893	1.	The first such call shall not set GLOB_APPEND. All subsequent calls shall set it.	
18894	2.	All the calls shall set GLOB_DOOFFS, or all shall not set it.	
18895	3.	After the second call, <i>pglob-&gt;gl_pathv</i> points to a list containing the following:	
18896	a.	Zero or more null pointers, as specified by GLOB_DOOFFS and <i>pglob-&gt;gl_offs</i> .	
18897	b.	Pointers to the pathnames that were in the <i>pglob-&gt;gl_pathv</i> list before the call, in the	
18898		same order as before.	
18899	c.	Pointers to the new pathnames generated by the second call, in the specified order.	
18900	4.	The count returned in <i>pglob-&gt;gl_pathc</i> shall be the total number of pathnames from the two	
18901		calls.	
18902	5.	The application can change any of the fields after a call to <i>glob()</i> . If it does, the application	
18903		shall reset them to the original value before a subsequent call, using the same <i>pglob</i> value,	
18904		to <i>globfree()</i> or <i>glob()</i> with the GLOB_APPEND flag.	
18905	If, during the search, a directory is encountered that cannot be opened or read and <i>errfunc</i> is not		
18906	a null pointer, <i>glob()</i> calls ( <i>*errfunc()</i> ) with two arguments:		
18907	1.	The <i>epath</i> argument is a pointer to the path that failed.	
18908	2.	The <i>errno</i> argument is the value of <i>errno</i> from the failure, as set by <i>opendir()</i> , <i>readdir()</i> , or	
18909		<i>stat()</i> . (Other values may be used to report other errors not explicitly documented for	
18910		those functions.)	
18911	If ( <i>*errfunc()</i> ) is called and returns non-zero, or if the GLOB_ERR flag is set in <i>flags</i> , <i>glob()</i> shall		
18912	stop the scan and return GLOB_ABORTED after setting <i>gl_pathc</i> and <i>gl_pathv</i> in <i>pglob</i> to reflect		
18913	the paths already scanned. If GLOB_ERR is not set and either <i>errfunc</i> is a null pointer or		
18914	<i>*errfunc()</i> returns 0, the error shall be ignored.		
18915	The <i>glob()</i> function shall not fail because of large files.		
18916	<b>RETURN VALUE</b>		
18917	Upon successful completion, <i>glob()</i> shall return 0. The argument <i>pglob-&gt;gl_pathc</i> shall return the		
18918	number of matched pathnames and the argument <i>pglob-&gt;gl_pathv</i> shall contain a pointer to a		
18919	null-terminated list of matched and sorted pathnames. However, if <i>pglob-&gt;gl_pathc</i> is 0, the		
18920	content of <i>pglob-&gt;gl_pathv</i> is undefined.		

18921 The *globfree()* function shall not return a value.

18922 If *glob()* terminates due to an error, it shall return one of the non-zero constants defined in  
18923 <glob.h>. The arguments *pglob->gl\_pathc* and *pglob->gl\_pathv* are still set as defined above.

#### 18924 ERRORS

18925 The *glob()* function shall fail and return the corresponding value if:

18926 GLOB\_ABORTED The scan was stopped because GLOB\_ERR was set or (*\*errfunc()*)  
18927 returned non-zero.

18928 GLOB\_NOMATCH The pattern does not match any existing pathname, and |  
18929 GLOB\_NOCHECK was not set in flags. |

18930 GLOB\_NOSPACE An attempt to allocate memory failed.

#### 18931 EXAMPLES

18932 One use of the GLOB\_DOOFFS flag is by applications that build an argument list for use with  
18933 *execv()*, *execve()*, or *execvp()*. Suppose, for example, that an application wants to do the  
18934 equivalent of:

```
18935 ls -l *.c
```

18936 but for some reason:

```
18937 system("ls -l *.c")
```

18938 is not acceptable. The application could obtain approximately the same result using the  
18939 sequence:

```
18940 globbuf.gl_offs = 2;  
18941 glob("*.c", GLOB_DOOFFS, NULL, &globbuf);  
18942 globbuf.gl_pathv[0] = "ls";  
18943 globbuf.gl_pathv[1] = "-l";  
18944 execvp("ls", &globbuf.gl_pathv[0]);
```

18945 Using the same example:

```
18946 ls -l *.c *.h
```

18947 could be approximately simulated using GLOB\_APPEND as follows:

```
18948 globbuf.gl_offs = 2;  
18949 glob("*.c", GLOB_DOOFFS, NULL, &globbuf);  
18950 glob("*.h", GLOB_DOOFFS|GLOB_APPEND, NULL, &globbuf);  
18951 ...
```

#### 18952 APPLICATION USAGE

18953 This function is not provided for the purpose of enabling utilities to perform pathname |  
18954 expansion on their arguments, as this operation is performed by the shell, and utilities are |  
18955 explicitly not expected to redo this. Instead, it is provided for applications that need to do |  
18956 pathname expansion on strings obtained from other sources, such as a pattern typed by a user or |  
18957 read from a file.

18958 If a utility needs to see if a pathname matches a given pattern, it can use *fnmatch()*. |

18959 Note that *gl\_pathc* and *gl\_pathv* have meaning even if *glob()* fails. This allows *glob()* to report  
18960 partial results in the event of an error. However, if *gl\_pathc* is 0, *gl\_pathv* is unspecified even if  
18961 *glob()* did not return an error.

18962 The GLOB\_NOCHECK option could be used when an application wants to expand a pathname |  
18963 if wildcards are specified, but wants to treat the pattern as just a string otherwise. The *sh* utility |

18964 might use this for option-arguments, for example.

18965 The new pathnames generated by a subsequent call with GLOB\_APPEND are not sorted |  
18966 together with the previous pathnames. This mirrors the way that the shell handles pathname |  
18967 expansion when multiple expansions are done on a command line. |

18968 Applications that need tilde and parameter expansion should use *wordexp()*.

#### 18969 RATIONALE

18970 It was claimed that the GLOB\_DOOFFS flag is unnecessary because it could be simulated using:

```
18971 new = (char **)malloc((n + pglob->gl_pathc + 1)
18972     * sizeof(char *));
18973 (void) memcpy(new+n, pglob->gl_pathv,
18974     pglob->gl_pathc * sizeof(char *));
18975 (void) memset(new, 0, n * sizeof(char *));
18976 free(pglob->gl_pathv);
18977 pglob->gl_pathv = new;
```

18978 However, this assumes that the memory pointed to by *gl\_pathv* is a block that was separately  
18979 created using *malloc()*. This is not necessarily the case. An application should make no  
18980 assumptions about how the memory referenced by fields in *pglob* was allocated. It might have  
18981 been obtained from *malloc()* in a large chunk and then carved up within *glob()*, or it might have  
18982 been created using a different memory allocator. It is not the intent of the standard developers to  
18983 specify or imply how the memory used by *glob()* is managed.

18984 The GLOB\_APPEND flag would be used when an application wants to expand several different  
18985 patterns into a single list.

#### 18986 FUTURE DIRECTIONS

18987 None.

#### 18988 SEE ALSO

18989 *exec*, *fnmatch()*, *opendir()*, *readdir()*, *stat()*, *wordexp()*, the Base Definitions volume of  
18990 IEEE Std 1003.1-200x, <**glob.h**>, the Shell and Utilities volume of IEEE Std 1003.1-200x

#### 18991 CHANGE HISTORY

18992 First released in Issue 4. Derived from the ISO POSIX-2 standard.

#### 18993 Issue 5

18994 Moved from POSIX2 C-language Binding to BASE.

#### 18995 Issue 6

18996 The DESCRIPTION is updated to avoid use of the term “must” for application requirements.

18997 The **restrict** keyword is added to the *glob()* prototype for alignment with the ISO/IEC 9899:1999  
18998 standard.



18999 **NAME**

19000 gmtime, gmtime\_r — convert a time value to a broken-down UTC time

19001 **SYNOPSIS**

19002 #include <time.h>

19003 struct tm \*gmtime(const time\_t \*timer);

19004 TSF struct tm \*gmtime\_r(const time\_t \*restrict timer,

19005 struct tm \*restrict result);

19006

19007 **DESCRIPTION**

19008 CX For *gmtime()*: The functionality described on this reference page is aligned with the ISO C  
 19009 standard. Any conflict between the requirements described here and the ISO C standard is  
 19010 unintentional. This volume of IEEE Std 1003.1-200x defers to the ISO C standard.

19011 The *gmtime()* function shall convert the time in seconds since the Epoch pointed to by *timer* into  
 19012 a broken-down time, expressed as Coordinated Universal Time (UTC).

19013 CX The relationship between a time in seconds since the Epoch used as an argument to *gmtime()* |  
 19014 and the **tm** structure (defined in the <time.h> header) is that the result shall be as specified in the |  
 19015 expression given in the definition of seconds since the Epoch (see the Base Definitions volume of |  
 19016 IEEE Std 1003.1-200x, Section 4.14, Seconds Since the Epoch), where the names in the structure |  
 19017 and in the expression correspond. |

19018 TSF The same relationship shall apply for *gmtime\_r()*. |

19019 CX The *gmtime()* function need not be reentrant. A function that is not required to be reentrant is not |  
 19020 required to be thread-safe. |

19021 The *asctime()*, *ctime()*, *gmtime()*, and *localtime()* functions shall return values in one of two static |  
 19022 objects: a broken-down time structure and an array of type **char**. Execution of any of the |  
 19023 functions may overwrite the information returned in either of these objects by any of the other |  
 19024 functions. |

19025 TSF The *gmtime\_r()* function shall convert the time in seconds since the Epoch pointed to by *timer* |  
 19026 into a broken-down time expressed as Coordinated Universal Time (UTC). The broken-down |  
 19027 time is stored in the structure referred to by *result*. The *gmtime\_r()* function shall also return the |  
 19028 address of the same structure. |

19029 **RETURN VALUE**

19030 The *gmtime()* function shall return a pointer to a **struct tm**.

19031 TSF Upon successful completion, *gmtime\_r()* shall return the address of the structure pointed to by |  
 19032 the argument *result*. If an error is detected, or UTC is not available, *gmtime\_r()* shall return a null |  
 19033 pointer. |

19034 **ERRORS**

19035 No errors are defined.

19036 **EXAMPLES**

19037 None.

19038 **APPLICATION USAGE**

19039 The *gmtime\_r()* function is thread-safe and returns values in a user-supplied buffer instead of  
19040 possibly using a static data area that may be overwritten by each call.

19041 **RATIONALE**

19042 None.

19043 **FUTURE DIRECTIONS**

19044 None.

19045 **SEE ALSO**

19046 *asctime()*, *clock()*, *ctime()*, *difftime()*, *localtime()*, *mktime()*, *strptime()*, *strptime()*, *time()*, *utime()*,  
19047 the Base Definitions volume of IEEE Std 1003.1-200x, <**time.h**>

19048 **CHANGE HISTORY**

19049 First released in Issue 1. Derived from Issue 1 of the SVID.

19050 **Issue 5**

19051 A note indicating that the *gmtime()* function need not be reentrant is added to the  
19052 DESCRIPTION.

19053 The *gmtime\_r()* function is included for alignment with the POSIX Threads Extension.19054 **Issue 6**19055 The *gmtime\_r()* function is marked as part of the Thread-Safe Functions option.

19056 Extensions beyond the ISO C standard are now marked.

19057 The APPLICATION USAGE section is updated to include a note on the thread-safe function and  
19058 its avoidance of possibly using a static data area.

19059 The **restrict** keyword is added to the *gmtime\_r()* prototype for alignment with the  
19060 ISO/IEC 9899:1999 standard.

19061 **NAME**

19062 grantpt — grant access to the slave pseudo-terminal device

19063 **SYNOPSIS**

19064 XSI #include &lt;stdlib.h&gt;

19065 int grantpt(int *fildev*);

19066

19067 **DESCRIPTION**

19068 The *grantpt()* function shall change the mode and ownership of the slave pseudo-terminal  
19069 device associated with its master pseudo-terminal counterpart. The *fildev* argument is a file  
19070 descriptor that refers to a master pseudo-terminal device. The user ID of the slave shall be set to  
19071 the real UID of the calling process and the group ID shall be set to an unspecified group ID. The  
19072 permission mode of the slave pseudo-terminal shall be set to readable and writable by the  
19073 owner, and writable by the group.

19074 The behavior of the *grantpt()* function is unspecified if the application has installed a signal  
19075 handler to catch SIGCHLD signals.

19076 **RETURN VALUE**

19077 Upon successful completion, *grantpt()* shall return 0; otherwise, it shall return -1 and set *errno* to  
19078 indicate the error.

19079 **ERRORS**19080 The *grantpt()* function may fail if:19081 [EBADF] The *fildev* argument is not a valid open file descriptor.19082 [EINVAL] The *fildev* argument is not associated with a master pseudo-terminal device.

19083 [EACCES] The corresponding slave pseudo-terminal device could not be accessed.

19084 **EXAMPLES**

19085 None.

19086 **APPLICATION USAGE**

19087 None.

19088 **RATIONALE**

19089 None.

19090 **FUTURE DIRECTIONS**

19091 None.

19092 **SEE ALSO**19093 *open()*, *ptsname()*, *unlockpt()*, the Base Definitions volume of IEEE Std 1003.1-200x, <stdlib.h>19094 **CHANGE HISTORY**

19095 First released in Issue 4, Version 2.

19096 **Issue 5**

19097 Moved from X/OPEN UNIX extension to BASE.

19098 The last paragraph of the DESCRIPTION is moved from the APPLICATION USAGE section in  
19099 previous issues.

19100 **NAME**

19101 h\_errno — error return value for network database operations

19102 **SYNOPSIS**

```
19103 OB #include <netdb.h>
```

19104

19105 **DESCRIPTION**

19106 Note that this method of returning errors is used only in connection with obsolescent functions.

19107 The <netdb.h> header provides a declaration of *h\_errno* as a modifiable *l*-value of type **int**.

19108 It is unspecified whether *h\_errno* is a macro or an identifier declared with external linkage. If a  
19109 macro definition is suppressed in order to access an actual object, or a program defines an  
19110 identifier with the name *h\_errno*, the behavior is undefined.

19111 **RETURN VALUE**

19112 None.

19113 **ERRORS**

19114 No errors are defined.

19115 **EXAMPLES**

19116 None.

19117 **APPLICATION USAGE**

19118 Applications should obtain the definition of *h\_errno* by the inclusion of the <netdb.h> header.

19119 **RATIONALE**

19120 None.

19121 **FUTURE DIRECTIONS**

19122 *h\_errno* may be withdrawn in a future version.

19123 **SEE ALSO**

19124 *endhostent()*, *errno*, the Base Definitions volume of IEEE Std 1003.1-200x, <netdb.h>

19125 **CHANGE HISTORY**

19126 First released in Issue 6. Derived from the XNS, Issue 5.2 specification.

19127 **NAME**

19128       hcreate, hdestroy, hsearch — manage hash search table

19129 **SYNOPSIS**

```

19130 xSI      #include <search.h>
19131          int hcreate(size_t nel);
19132          void hdestroy(void);
19133          ENTRY *hsearch(ENTRY item, ACTION action);
19134

```

19135 **DESCRIPTION**19136       The *hcreate()*, *hdestroy()*, and *hsearch()* functions shall manage hash search tables. |

19137       The *hcreate()* function shall allocate sufficient space for the table, and the application shall |  
 19138       ensure it is called before *hsearch()* is used. The *nel* argument is an estimate of the maximum |  
 19139       number of entries that the table shall contain. This number may be adjusted upward by the |  
 19140       algorithm in order to obtain certain mathematically favorable circumstances.

19141       The *hdestroy()* function shall dispose of the search table, and may be followed by another call to |  
 19142       *hcreate()*. After the call to *hdestroy()*, the data can no longer be considered accessible.

19143       The *hsearch()* function is a hash-table search routine. It shall return a pointer into a hash table |  
 19144       indicating the location at which an entry can be found. The *item* argument is a structure of type |  
 19145       **ENTRY** (defined in the *<search.h>* header) containing two pointers: *item.key* points to the |  
 19146       comparison key (a **char** \*), and *item.data* (a **void** \*) points to any other data to be associated with |  
 19147       that key. The comparison function used by *hsearch()* is *strcmp()*. The *action* argument is a |  
 19148       member of an enumeration type **ACTION** indicating the disposition of the entry if it cannot be |  
 19149       found in the table. **ENTER** indicates that the item should be inserted in the table at an |  
 19150       appropriate point. **FIND** indicates that no entry should be made. Unsuccessful resolution is |  
 19151       indicated by the return of a null pointer.

19152       These functions need not be reentrant. A function that is not required to be reentrant is not |  
 19153       required to be thread-safe.

19154 **RETURN VALUE**

19155       The *hcreate()* function shall return 0 if it cannot allocate sufficient space for the table; otherwise, |  
 19156       it shall return non-zero.

19157       The *hdestroy()* function shall not return a value.

19158       The *hsearch()* function shall return a null pointer if either the action is **FIND** and the item could |  
 19159       not be found or the action is **ENTER** and the table is full.

19160 **ERRORS**

19161       The *hcreate()* and *hsearch()* functions may fail if:

19162       [ENOMEM]       Insufficient storage space is available.

19163 **EXAMPLES**

19164       The following example reads in strings followed by two numbers and stores them in a hash |  
 19165       table, discarding duplicates. It then reads in strings and finds the matching entry in the hash |  
 19166       table and prints it out.

```

19167          #include <stdio.h>
19168          #include <search.h>
19169          #include <string.h>
19170          struct info {          /* This is the info stored in the table */
19171              int age, room;     /* other than the key. */

```

```

19172     };
19173     #define NUM_EMPL    5000    /* # of elements in search table. */
19174     int main(void)
19175     {
19176         char string_space[NUM_EMPL*20];    /* Space to store strings. */
19177         struct info info_space[NUM_EMPL]; /* Space to store employee info. */
19178         char *str_ptr = string_space;     /* Next space in string_space. */
19179         struct info *info_ptr = info_space;
19180                                         /* Next space in info_space. */
19181         ENTRY item;
19182         ENTRY *found_item; /* Name to look for in table. */
19183         char name_to_find[30];
19184
19184         int i = 0;
19185
19185         /* Create table; no error checking is performed. */
19186         (void) hcreate(NUM_EMPL);
19187         while (scanf("%s%d%d", str_ptr, &info_ptr->age,
19188             &info_ptr->room) != EOF && i++ < NUM_EMPL) {
19189
19189             /* Put information in structure, and structure in item. */
19190             item.key = str_ptr;
19191             item.data = info_ptr;
19192             str_ptr += strlen(str_ptr) + 1;
19193             info_ptr++;
19194
19194             /* Put item into table. */
19195             (void) hsearch(item, ENTER);
19196         }
19197
19197         /* Access table. */
19198         item.key = name_to_find;
19199         while (scanf("%s", item.key) != EOF) {
19200             if ((found_item = hsearch(item, FIND)) != NULL) {
19201
19201                 /* If item is in the table. */
19202                 (void)printf("found %s, age = %d, room = %d\n",
19203                     found_item->key,
19204                     ((struct info *)found_item->data)->age,
19205                     ((struct info *)found_item->data)->room);
19206             } else
19207                 (void)printf("no such employee %s\n", name_to_find);
19208         }
19209         return 0;
19210     }

```

#### 19211 APPLICATION USAGE

19212 The *hcreate()* and *hsearch()* functions may use *malloc()* to allocate space.

#### 19213 RATIONALE

19214 None.

19215 **FUTURE DIRECTIONS**

19216       None.

19217 **SEE ALSO**19218       *bsearch()*, *lsearch()*, *malloc()*, *strcmp()*, *tsearch()*, the Base Definitions volume of  
19219       IEEE Std 1003.1-200x, <**search.h**>19220 **CHANGE HISTORY**

19221       First released in Issue 1. Derived from Issue 1 of the SVID.

19222 **Issue 6**

19223       The DESCRIPTION is updated to avoid use of the term “must” for application requirements.

19224       A note indicating that this function need not be reentrant is added to the DESCRIPTION.

19225 **NAME**

19226 hdestroy — manage hash search table

19227 **SYNOPSIS**

19228 XSI #include <search.h>

19229 void hdestroy(void);

19230

19231 **DESCRIPTION**

19232 Refer to *hcreate()*.



19233 **NAME**

19234 hsearch — manage hash search table

19235 **SYNOPSIS**

19236 xSI #include &lt;search.h&gt;

19237 ENTRY \*hsearch(ENTRY *item*, ACTION *action*);

19238

19239 **DESCRIPTION**19240 Refer to *hcreate()*.

19241 **NAME**

19242 htonl, htons, ntohl, ntohs — convert values between host and network byte order

19243 **SYNOPSIS**

19244 #include <arpa/inet.h>

19245 uint32\_t htonl(uint32\_t *hostlong*);

19246 uint16\_t htons(uint16\_t *hostshort*);

19247 uint32\_t ntohl(uint32\_t *netlong*);

19248 uint16\_t ntohs(uint16\_t *netshort*);

19249 **DESCRIPTION**

19250 These functions shall convert 16-bit and 32-bit quantities between network byte order and host  
19251 byte order.

19252 On some implementations, these functions are defined as macros.

19253 The **uint32\_t** and **uint16\_t** types are defined in <inttypes.h>.

19254 **RETURN VALUE**

19255 The *htonl()* and *htons()* functions shall return the argument value converted from host to  
19256 network byte order.

19257 The *ntohl()* and *ntohs()* functions shall return the argument value converted from network to  
19258 host byte order.

19259 **ERRORS**

19260 No errors are defined.

19261 **EXAMPLES**

19262 None.

19263 **APPLICATION USAGE**

19264 These functions are most often used in conjunction with IPv4 addresses and ports as returned by  
19265 *gethostent()* and *getservent()*.

19266 **RATIONALE**

19267 None.

19268 **FUTURE DIRECTIONS**

19269 None.

19270 **SEE ALSO**

19271 *endhostent()*, *endservent()*, the Base Definitions volume of IEEE Std 1003.1-200x, <inttypes.h>,  
19272 <arpa/inet.h>

19273 **CHANGE HISTORY**

19274 First released in Issue 6. Derived from the XNS, Issue 5.2 specification.

19275 **NAME**

19276       htons — convert values between host and network byte order

19277 **SYNOPSIS**

19278       #include <arpa/inet.h>

19279       uint16\_t htons(uint16\_t *hostshort*);

19280 **DESCRIPTION**

19281       Refer to *htonl()*.

19282 **NAME**

19283 hypot, hypotf, hypotl — Euclidean distance function

19284 **SYNOPSIS**

19285 #include &lt;math.h&gt;

19286 double hypot(double x, double y);

19287 float hypotf(float x, float y);

19288 long double hypotl(long double x, long double y);

19289 **DESCRIPTION**

19290 **CX** The functionality described on this reference page is aligned with the ISO C standard. Any  
 19291 conflict between the requirements described here and the ISO C standard is unintentional. This  
 19292 volume of IEEE Std 1003.1-200x defers to the ISO C standard.

19293 These functions shall compute the value of the square root of  $x^2+y^2$  without undue overflow or  
 19294 underflow.

19295 An application wishing to check for error situations should set *errno* to zero and call  
 19296 *feclearexcept*(FE\_ALL\_EXCEPT) before calling these functions. On return, if *errno* is non-zero or  
 19297 *fetestexcept*(FE\_INVALID | FE\_DIVBYZERO | FE\_OVERFLOW | FE\_UNDERFLOW) is non-  
 19298 zero, an error has occurred.

19299 **RETURN VALUE**

19300 Upon successful completion, these functions shall return the length of the hypotenuse of a  
 19301 right-angled triangle with sides of length *x* and *y*.

19302 If the correct value would cause overflow, a range error shall occur and *hypot()*, *hypotf()*, and  
 19303 *hypotl()* shall return the value of the macro HUGE\_VAL, HUGE\_VALF, and HUGE\_VALL,  
 19304 respectively.

19305 **MX** If *x* or *y* is  $\pm\text{Inf}$ ,  $+\text{Inf}$  shall be returned (even if one of *x* or *y* is NaN).

19306 If *x* or *y* is NaN, and the other is not  $\pm\text{Inf}$ , a NaN shall be returned.

19307 If both arguments are subnormal and the correct result is subnormal, a range error may occur  
 19308 and the correct result is returned.

19309 **ERRORS**

19310 These functions shall fail if:

19311 **Range Error** The result overflows.

19312 If the integer expression (math\_errhandling & MATH\_ERRNO) is non-zero, |  
 19313 then *errno* shall be set to [ERANGE]. If the integer expression |  
 19314 (math\_errhandling & MATH\_ERREXCEPT) is non-zero, then the overflow |  
 19315 floating-point exception shall be raised. |

19316 These functions may fail if:

19317 **MX** **Range Error** The result underflows.

19318 If the integer expression (math\_errhandling & MATH\_ERRNO) is non-zero, |  
 19319 then *errno* shall be set to [ERANGE]. If the integer expression |  
 19320 (math\_errhandling & MATH\_ERREXCEPT) is non-zero, then the underflow |  
 19321 floating-point exception shall be raised. |

19322 **EXAMPLES**

19323 None.

19324 **APPLICATION USAGE**19325 *hypot(x,y)*, *hypot(y,x)*, and *hypot(x,-y)* are equivalent.19326 *hypot(x, ±0)* is equivalent to *fabs(x)*.19327 Underflow only happens when both *x* and *y* are subnormal and the (inexact) result is also  
19328 subnormal.19329 These functions take precautions against overflow during intermediate steps of the  
19330 computation.19331 On error, the expressions (math\_errhandling & MATH\_ERRNO) and (math\_errhandling &  
19332 MATH\_ERREXCEPT) are independent of each other, but at least one of them must be non-zero.19333 **RATIONALE**

19334 None.

19335 **FUTURE DIRECTIONS**

19336 None.

19337 **SEE ALSO**19338 *feclearexcept()*, *fetestexcept()*, *isnan()*, *sqrt()*, the Base Definitions volume of IEEE Std 1003.1-200x, |  
19339 Section 4.18, Treatment of Error Conditions for Mathematical Functions, <math.h> |19340 **CHANGE HISTORY**

19341 First released in Issue 1. Derived from Issue 1 of the SVID.

19342 **Issue 5**19343 The DESCRIPTION is updated to indicate how an application should check for an error. This  
19344 text was previously published in the APPLICATION USAGE section.19345 **Issue 6**19346 The *hypot()* function is no longer marked as an extension.19347 The *hypotf()* and *hypotl()* functions are added for alignment with the ISO/IEC 9899:1999  
19348 standard.19349 The DESCRIPTION, RETURN VALUE, ERRORS, and APPLICATION USAGE sections are  
19350 revised to align with the ISO/IEC 9899:1999 standard.19351 IEC 60559:1989 standard floating-point extensions over the ISO/IEC 9899:1999 standard are  
19352 marked.

## 19353 NAME

19354 iconv — codeset conversion function

## 19355 SYNOPSIS

19356 XSI `#include <iconv.h>`

```
19357     size_t iconv(iconv_t cd, char **restrict inbuf,
19358                size_t *restrict inbytesleft, char **restrict outbuf,
19359                size_t *restrict outbytesleft);
19360
```

## 19361 DESCRIPTION

19362 The *iconv()* function shall convert the sequence of characters from one codeset, in the array  
 19363 specified by *inbuf*, into a sequence of corresponding characters in another codeset, in the array  
 19364 specified by *outbuf*. The codesets are those specified in the *iconv\_open()* call that returned the  
 19365 conversion descriptor, *cd*. The *inbuf* argument points to a variable that points to the first  
 19366 character in the input buffer and *inbytesleft* indicates the number of bytes to the end of the buffer  
 19367 to be converted. The *outbuf* argument points to a variable that points to the first available byte in  
 19368 the output buffer and *outbytesleft* indicates the number of the available bytes to the end of the  
 19369 buffer.

19370 For state-dependent encodings, the conversion descriptor *cd* is placed into its initial shift state by  
 19371 a call for which *inbuf* is a null pointer, or for which *inbuf* points to a null pointer. When *iconv()* is  
 19372 called in this way, and if *outbuf* is not a null pointer or a pointer to a null pointer, and *outbytesleft*  
 19373 points to a positive value, *iconv()* shall place, into the output buffer, the byte sequence to change  
 19374 the output buffer to its initial shift state. If the output buffer is not large enough to hold the  
 19375 entire reset sequence, *iconv()* shall fail and set *errno* to [E2BIG]. Subsequent calls with *inbuf* as  
 19376 other than a null pointer or a pointer to a null pointer cause the conversion to take place from  
 19377 the current state of the conversion descriptor.

19378 If a sequence of input bytes does not form a valid character in the specified codeset, conversion |  
 19379 shall stop after the previous successfully converted character. If the input buffer ends with an |  
 19380 incomplete character or shift sequence, conversion shall stop after the previous successfully |  
 19381 converted bytes. If the output buffer is not large enough to hold the entire converted input, |  
 19382 conversion shall stop just prior to the input bytes that would cause the output buffer to |  
 19383 overflow. The variable pointed to by *inbuf* shall be updated to point to the byte following the last |  
 19384 byte successfully used in the conversion. The value pointed to by *inbytesleft* shall be |  
 19385 decremented to reflect the number of bytes still not converted in the input buffer. The variable |  
 19386 pointed to by *outbuf* shall be updated to point to the byte following the last byte of converted |  
 19387 output data. The value pointed to by *outbytesleft* shall be decremented to reflect the number of |  
 19388 bytes still available in the output buffer. For state-dependent encodings, the conversion |  
 19389 descriptor shall be updated to reflect the shift state in effect at the end of the last successfully |  
 19390 converted byte sequence.

19391 If *iconv()* encounters a character in the input buffer that is valid, but for which an identical |  
 19392 character does not exist in the target codeset, *iconv()* shall perform an implementation-defined |  
 19393 conversion on this character. |

## 19394 RETURN VALUE

19395 The *iconv()* function shall update the variables pointed to by the arguments to reflect the extent  
 19396 of the conversion and return the number of non-identical conversions performed. If the entire  
 19397 string in the input buffer is converted, the value pointed to by *inbytesleft* shall be 0. If the input  
 19398 conversion is stopped due to any conditions mentioned above, the value pointed to by *inbytesleft*  
 19399 shall be non-zero and *errno* shall be set to indicate the condition. If an error occurs *iconv()* shall  
 19400 return (*size\_t*)-1 and set *errno* to indicate the error.

19401 **ERRORS**

19402 The *iconv()* function shall fail if:

19403 [EILSEQ] Input conversion stopped due to an input byte that does not belong to the  
19404 input codeset.

19405 [E2BIG] Input conversion stopped due to lack of space in the output buffer.

19406 [EINVAL] Input conversion stopped due to an incomplete character or shift sequence at  
19407 the end of the input buffer.

19408 The *iconv()* function may fail if:

19409 [EBADF] The *cd* argument is not a valid open conversion descriptor.

19410 **EXAMPLES**

19411 None.

19412 **APPLICATION USAGE**

19413 The *inbuf* argument indirectly points to the memory area which contains the conversion input  
19414 data. The *outbuf* argument indirectly points to the memory area which is to contain the result of  
19415 the conversion. The objects indirectly pointed to by *inbuf* and *outbuf* are not restricted to  
19416 containing data that is directly representable in the ISO C standard language **char** data type. The  
19417 type of *inbuf* and *outbuf*, **char \*\***, does not imply that the objects pointed to are interpreted as  
19418 null-terminated C strings or arrays of characters. Any interpretation of a byte sequence that  
19419 represents a character in a given character set encoding scheme is done internally within the  
19420 codeset converters. For example, the area pointed to indirectly by *inbuf* and/or *outbuf* can  
19421 contain all zero octets that are not interpreted as string terminators but as coded character data  
19422 according to the respective codeset encoding scheme. The type of the data (**char**, **short**, **long**, and  
19423 so on) read or stored in the objects is not specified, but may be inferred for both the input and  
19424 output data by the converters determined by the *fromcode* and *toctype* arguments of *iconv\_open()*.

19425 Regardless of the data type inferred by the converter, the size of the remaining space in both  
19426 input and output objects (the *inbytesleft* and *outbytesleft* arguments) is always measured in bytes.

19427 For implementations that support the conversion of state-dependent encodings, the conversion  
19428 descriptor must be able to accurately reflect the shift-state in effect at the end of the last  
19429 successful conversion. It is not required that the conversion descriptor itself be updated, which  
19430 would require it to be a pointer type. Thus, implementations are free to implement the  
19431 descriptor as a handle (other than a pointer type) by which the conversion information can be  
19432 accessed and updated.

19433 **RATIONALE**

19434 None.

19435 **FUTURE DIRECTIONS**

19436 None.

19437 **SEE ALSO**

19438 *iconv\_open()*, *iconv\_close()*, the Base Definitions volume of IEEE Std 1003.1-200x, <**iconv.h**>

19439 **CHANGE HISTORY**

19440 First released in Issue 4. Derived from the HP-UX Manual.

19441 **Issue 6**

19442 The SYNOPSIS has been corrected to align with the <**iconv.h**> reference page.

19443 The **restrict** keyword is added to the *iconv()* prototype for alignment with the  
19444 ISO/IEC 9899:1999 standard.

19445 **NAME**

19446 iconv\_close — codeset conversion deallocation function

19447 **SYNOPSIS**

19448 XSI #include &lt;iconv.h&gt;

19449 int iconv\_close(iconv\_t cd);

19450

19451 **DESCRIPTION**19452 The *iconv\_close()* function shall deallocate the conversion descriptor *cd* and all other associated |  
19453 resources allocated by *iconv\_open()*.19454 If a file descriptor is used to implement the type **iconv\_t**, that file descriptor shall be closed. |19455 **RETURN VALUE**19456 Upon successful completion, 0 shall be returned; otherwise, -1 shall be returned and *errno* set to |  
19457 indicate the error.19458 **ERRORS**19459 The *iconv\_close()* function may fail if:

19460 [EBADF] The conversion descriptor is invalid.

19461 **EXAMPLES**

19462 None.

19463 **APPLICATION USAGE**

19464 None.

19465 **RATIONALE**

19466 None.

19467 **FUTURE DIRECTIONS**

19468 None.

19469 **SEE ALSO**19470 *iconv()*, *iconv\_open()*, the Base Definitions volume of IEEE Std 1003.1-200x, <**iconv.h**>19471 **CHANGE HISTORY**

19472 First released in Issue 4. Derived from the HP-UX Manual.



19473 **NAME**

19474 iconv\_open — codeset conversion allocation function

19475 **SYNOPSIS**19476 XSI `#include <iconv.h>`19477 `iconv_t iconv_open(const char *tocode, const char *fromcode);`

19478

19479 **DESCRIPTION**

19480 The *iconv\_open()* function shall return a conversion descriptor that describes a conversion from  
 19481 the codeset specified by the string pointed to by the *fromcode* argument to the codeset specified  
 19482 by the string pointed to by the *tocode* argument. For state-dependent encodings, the conversion  
 19483 descriptor shall be in a codeset-dependent initial shift state, ready for immediate use with  
 19484 *iconv()*.

19485 Settings of *fromcode* and *tocode* and their permitted combinations are implementation-defined.

19486 A conversion descriptor shall remain valid until it is closed by *iconv\_close()* or an implicit close.

19487 If a file descriptor is used to implement conversion descriptors, the FD\_CLOEXEC flag shall be  
 19488 set; see <fcntl.h>.

19489 **RETURN VALUE**

19490 Upon successful completion, *iconv\_open()* shall return a conversion descriptor for use on  
 19491 subsequent calls to *iconv()*. Otherwise, *iconv\_open()* shall return (**iconv\_t**)-1 and set *errno* to  
 19492 indicate the error.

19493 **ERRORS**

19494 The *iconv\_open()* function may fail if:

19495 [EMFILE] {OPEN\_MAX} file descriptors are currently open in the calling process.

19496 [ENFILE] Too many files are currently open in the system.

19497 [ENOMEM] Insufficient storage space is available.

19498 [EINVAL] The conversion specified by *fromcode* and *tocode* is not supported by the  
 19499 implementation.

19500 **EXAMPLES**

19501 None.

19502 **APPLICATION USAGE**

19503 Some implementations of *iconv\_open()* use *malloc()* to allocate space for internal buffer areas.  
 19504 The *iconv\_open()* function may fail if there is insufficient storage space to accommodate these  
 19505 buffers.

19506 Conforming applications must assume that conversion descriptors are not valid after a call to  
 19507 one of the *exec* functions.

19508 **RATIONALE**

19509 None.

19510 **FUTURE DIRECTIONS**

19511 None.

19512 **SEE ALSO**

19513 *iconv()*, *iconv\_close()*, the Base Definitions volume of IEEE Std 1003.1-200x, <fcntl.h>, <iconv.h>

19514 **CHANGE HISTORY**

19515 First released in Issue 4. Derived from the HP-UX Manual.

19516 **NAME**

19517 if\_freenameindex — free memory allocated by *if\_nameindex()*

19518 **SYNOPSIS**

19519 #include <net/if.h>

19520 void if\_freenameindex(struct if\_nameindex \*ptr);

19521 **DESCRIPTION**

19522 The *if\_freenameindex()* function shall free the memory allocated by *if\_nameindex()*. The *ptr*  
19523 argument shall be a pointer that was returned by *if\_nameindex()*. After *if\_freenameindex()* has  
19524 been called, the application shall not use the array of which *ptr* is the address. |

19525 **RETURN VALUE**

19526 None.

19527 **ERRORS**

19528 No errors are defined.

19529 **EXAMPLES**

19530 None.

19531 **APPLICATION USAGE**

19532 None.

19533 **RATIONALE**

19534 None.

19535 **FUTURE DIRECTIONS**

19536 None.

19537 **SEE ALSO**

19538 *getsockopt()*, *if\_indextoname()*, *if\_nameindex()*, *if\_nametoindex()*, *setsockopt()*, the Base Definitions  
19539 volume of IEEE Std 1003.1-200x, <net/if.h>

19540 **CHANGE HISTORY**

19541 First released in Issue 6. Derived from the XNS, Issue 5.2 specification.

19542 **NAME**

19543 if\_indextoname — map a network interface index to its corresponding name

19544 **SYNOPSIS**

19545 #include <net/if.h>

19546 char \*if\_indextoname(unsigned *ifindex*, char \**ifname*);

19547 **DESCRIPTION**

19548 The *if\_indextoname*() function shall map an interface index to its corresponding name.

19549 When this function is called, *ifname* shall point to a buffer of at least {IFNAMSIZ} bytes. The  
19550 function shall place in this buffer the name of the interface with index *ifindex*.

19551 **RETURN VALUE**

19552 If *ifindex* is an interface index, then the function shall return the value supplied in *ifname*, which  
19553 points to a buffer now containing the interface name. Otherwise, the function shall return a  
19554 NULL pointer and set *errno* to indicate the error.

19555 **ERRORS**

19556 The *if\_indextoname*() function shall fail if:

19557 [ENXIO] The interface does not exist.

19558 **EXAMPLES**

19559 None.

19560 **APPLICATION USAGE**

19561 None.

19562 **RATIONALE**

19563 None.

19564 **FUTURE DIRECTIONS**

19565 None.

19566 **SEE ALSO**

19567 *getsockopt*(), *if\_freenameindex*(), *if\_nameindex*(), *if\_nametoindex*(), *setsockopt*(), the Base  
19568 Definitions volume of IEEE Std 1003.1-200x, <net/if.h>

19569 **CHANGE HISTORY**

19570 First released in Issue 6. Derived from the XNS, Issue 5.2 specification.

19571 **NAME**

19572 if\_nameindex — return all network interface names and indexes

19573 **SYNOPSIS**

19574 #include <net/if.h>

19575 struct if\_nameindex \*if\_nameindex(void);

19576 **DESCRIPTION**

19577 The *if\_nameindex()* function shall return an array of *if\_nameindex* structures, one structure per  
19578 interface. The end of the array is indicated by a structure with an *if\_index* field of zero and an  
19579 *if\_name* field of NULL.

19580 Applications should call *if\_freenameindex()* to release the memory that may be dynamically  
19581 allocated by this function, after they have finished using it.

19582 **RETURN VALUE**

19583 Array of structures identifying local interfaces. A NULL pointer is returned upon an error, with  
19584 *errno* set to indicate the error.

19585 **ERRORS**

19586 The *if\_nameindex()* function may fail if:

19587 [ENOBUFS] Insufficient resources are available to complete the function.

19588 **EXAMPLES**

19589 None.

19590 **APPLICATION USAGE**

19591 None.

19592 **RATIONALE**

19593 None.

19594 **FUTURE DIRECTIONS**

19595 None.

19596 **SEE ALSO**

19597 *getsockopt()*, *if\_freenameindex()*, *if\_indextoname()*, *if\_nametoindex()*, *setsockopt()*, the Base  
19598 Definitions volume of IEEE Std 1003.1-200x, <net/if.h>

19599 **CHANGE HISTORY**

19600 First released in Issue 6. Derived from the XNS, Issue 5.2 specification.

19601 **NAME**

19602       if\_nametoindex — map a network interface name to its corresponding index

19603 **SYNOPSIS**

19604       #include <net/if.h>

19605       unsigned if\_nametoindex(const char \*ifname);

19606 **DESCRIPTION**

19607       The *if\_nametoindex()* function shall return the interface index corresponding to name *ifname*.

19608 **RETURN VALUE**

19609       The corresponding index if *ifname* is the name of an interface; otherwise, zero.

19610 **ERRORS**

19611       No errors are defined.

19612 **EXAMPLES**

19613       None.

19614 **APPLICATION USAGE**

19615       None.

19616 **RATIONALE**

19617       None.

19618 **FUTURE DIRECTIONS**

19619       None.

19620 **SEE ALSO**

19621       *getsockopt()*, *if\_freenameindex()*, *if\_indextoname()*, *if\_nameindex()*, *setsockopt()*, the Base  
19622       Definitions volume of IEEE Std 1003.1-200x, <net/if.h>

19623 **CHANGE HISTORY**

19624       First released in Issue 6. Derived from the XNS, Issue 5.2 specification.

19625 **NAME**

19626           ilogb, ilogbf, ilogbl — return an unbiased exponent

19627 **SYNOPSIS**

```
19628           #include <math.h>
19629           int ilogb(double x);
19630           int ilogbf(float x);
19631           int ilogbl(long double x);
```

19632 **DESCRIPTION**

19633 **CX**       The functionality described on this reference page is aligned with the ISO C standard. Any  
 19634 conflict between the requirements described here and the ISO C standard is unintentional. This  
 19635 volume of IEEE Std 1003.1-200x defers to the ISO C standard.

19636       These functions shall return the exponent part of their argument *x*. Formally, the return value is  
 19637 the integral part of  $\log_r |x|$  as a signed integral value, for non-zero *x*, where *r* is the radix of the  
 19638 machine's floating-point arithmetic, which is the value of FLT\_RADIX defined in <float.h>.

19639       An application wishing to check for error situations should set *errno* to zero and call  
 19640 *feclearexcept*(FE\_ALL\_EXCEPT) before calling these functions. On return, if *errno* is non-zero or  
 19641 *fetestexcept*(FE\_INVALID | FE\_DIVBYZERO | FE\_OVERFLOW | FE\_UNDERFLOW) is non-  
 19642 zero, an error has occurred.

19643 **RETURN VALUE**

19644       Upon successful completion, these functions shall return the exponent part of *x* as a signed  
 19645 integer value. They are equivalent to calling the corresponding *logb*() function and casting the  
 19646 returned value to type **int**.

19647 **XSI**       If *x* is 0, a domain error shall occur, and the value FP\_ILOGB0 shall be returned.

19648 **XSI**       If *x* is  $\pm\text{Inf}$ , a domain error shall occur, and the value {INT\_MAX} shall be returned.

19649 **XSI**       If *x* is a NaN, a domain error shall occur, and the value FP\_ILOGBNAN shall be returned.

19650 **XSI**       If the correct value is greater than {INT\_MAX}, {INT\_MAX} shall be returned and a domain error  
 19651 shall occur.

19652       If the correct value is less than {INT\_MIN}, {INT\_MIN} shall be returned and a domain error  
 19653 shall occur.

19654 **ERRORS**

19655       These functions shall fail if:

19656 <b>XSI</b>	<b>Domain Error</b>	The <i>x</i> argument is zero, NaN, or $\pm\text{Inf}$ , or the correct value is not representable as an integer.
------------------	---------------------	---

19658	If the integer expression (math_errhandling & MATH_ERRNO) is non-zero,	
19659	then <i>errno</i> shall be set to [EDOM]. If the integer expression (math_errhandling	
19660	& MATH_ERREXCEPT) is non-zero, then the invalid floating-point exception	
19661	shall be raised.	

19662 **EXAMPLES**

19663           None.

19664 **APPLICATION USAGE**

19665           On error, the expressions (math\_errhandling & MATH\_ERRNO) and (math\_errhandling &  
19666           MATH\_ERREXCEPT) are independent of each other, but at least one of them must be non-zero.

19667 **RATIONALE**

19668           The errors come from taking the expected floating-point value and converting it to **int**, which is  
19669           invalid operation in IEEE Std 754-1985 (since overflow, infinity, and NaN are not representable  
19670           in a type **int**), so should be a domain error.

19671           There are no known implementations that overflow. For overflow to happen, {INT\_MAX} must  
19672           be less than  $LDBL\_MAX\_EXP * \log_2(\text{FLT\_RADIX})$  or {INT\_MIN} must be greater than  
19673            $LDBL\_MIN\_EXP * \log_2(\text{FLT\_RADIX})$  if subnormals are not supported, or {INT\_MIN} must be  
19674           greater than  $(LDBL\_MIN\_EXP - LDBL\_MANT\_DIG) * \log_2(\text{FLT\_RADIX})$  if subnormals are  
19675           supported.

19676 **FUTURE DIRECTIONS**

19677           None.

19678 **SEE ALSO**

19679           *feclearexcept()*, *fetestexcept()*, *logb()*, *scalb()*, the Base Definitions volume of IEEE Std 1003.1-200x, |  
19680           Section 4.18, Treatment of Error Conditions for Mathematical Functions, <float.h>, <math.h> |

19681 **CHANGE HISTORY**

19682           First released in Issue 4, Version 2.

19683 **Issue 5**

19684           Moved from X/OPEN UNIX extension to BASE.

19685 **Issue 6**19686           The *ilogb()* function is no longer marked as an extension.

19687           The *ilogbf()* and *ilogbl()* functions are added for alignment with the ISO/IEC 9899:1999  
19688           standard.

19689           The RETURN VALUE section is revised for alignment with the ISO/IEC 9899:1999 standard.

19690           XSI extensions are marked.



19691 **NAME**

19692           imaxabs — return absolute value

19693 **SYNOPSIS**

19694           #include &lt;inttypes.h&gt;

19695           intmax\_t imaxabs(intmax\_t j);

19696 **DESCRIPTION**

19697 cx       The functionality described on this reference page is aligned with the ISO C standard. Any  
19698       conflict between the requirements described here and the ISO C standard is unintentional. This  
19699       volume of IEEE Std 1003.1-200x defers to the ISO C standard.

19700       The *imaxabs()* function shall compute the absolute value of an integer *j*. If the result cannot be  
19701       represented, the behavior is undefined.

19702 **RETURN VALUE**19703       The *imaxabs()* function shall return the absolute value.19704 **ERRORS**

19705       No errors are defined.

19706 **EXAMPLES**

19707       None.

19708 **APPLICATION USAGE**

19709       The absolute value of the most negative number cannot be represented in two's complement.

19710 **RATIONALE**

19711       None.

19712 **FUTURE DIRECTIONS**

19713       None.

19714 **SEE ALSO**19715       *imaxdiv()*, the Base Definitions volume of IEEE Std 1003.1-200x, <inttypes.h>19716 **CHANGE HISTORY**

19717       First released in Issue 6. Derived from the ISO/IEC 9899:1999 standard.

19718 **NAME**

19719 imaxdiv — return quotient and remainder

19720 **SYNOPSIS**

19721 #include <inttypes.h>

19722 imaxdiv\_t imaxdiv(intmax\_t numer, intmax\_t denom);

19723 **DESCRIPTION**

19724 cx The functionality described on this reference page is aligned with the ISO C standard. Any  
19725 conflict between the requirements described here and the ISO C standard is unintentional. This  
19726 volume of IEEE Std 1003.1-200x defers to the ISO C standard.

19727 The *imaxdiv()* function shall compute *numer* / *denom* and *numer* % *denom* in a single operation.

19728 **RETURN VALUE**

19729 The *imaxdiv()* function shall return a structure of type **imaxdiv\_t**, comprising both the quotient  
19730 and the remainder. The structure shall contain (in either order) the members *quot* (the quotient)  
19731 and *rem* (the remainder), each of which has type **intmax\_t**.

19732 If either part of the result cannot be represented, the behavior is undefined.

19733 **ERRORS**

19734 No errors are defined.

19735 **EXAMPLES**

19736 None.

19737 **APPLICATION USAGE**

19738 None.

19739 **RATIONALE**

19740 None.

19741 **FUTURE DIRECTIONS**

19742 None.

19743 **SEE ALSO**

19744 *imaxabs()*, the Base Definitions volume of IEEE Std 1003.1-200x, <inttypes.h>

19745 **CHANGE HISTORY**

19746 First released in Issue 6. Derived from the ISO/IEC 9899:1999 standard.

19747 **NAME**19748 index — character string operations (**LEGACY**)19749 **SYNOPSIS**19750 XSI `#include <strings.h>`19751 `char *index(const char *s, int c);`

19752

19753 **DESCRIPTION**19754 The *index()* function shall be equivalent to *strchr()*.19755 **RETURN VALUE**19756 See *strchr()*.19757 **ERRORS**19758 See *strchr()*.19759 **EXAMPLES**

19760 None.

19761 **APPLICATION USAGE**19762 *strchr()* is preferred over this function.19763 For maximum portability, it is recommended to replace the function call to *index()* as follows:19764 `#define index(a,b) strchr((a),(b))`19765 **RATIONALE**

19766 None.

19767 **FUTURE DIRECTIONS**

19768 This function may be withdrawn in a future version.

19769 **SEE ALSO**19770 *strchr()*, the Base Definitions volume of IEEE Std 1003.1-200x, **<strings.h>**19771 **CHANGE HISTORY**

19772 First released in Issue 4, Version 2.

19773 **Issue 5**

19774 Moved from X/OPEN UNIX extension to BASE.

19775 **Issue 6**

19776 This function is marked LEGACY.

19777 **NAME**

19778 inet\_addr, inet\_ntoa — IPv4 address manipulation

19779 **SYNOPSIS**

19780 #include &lt;arpa/inet.h&gt;

19781 inet\_addr\_t inet\_addr(const char \*cp);

19782 char \*inet\_ntoa(struct in\_addr in);

19783 **DESCRIPTION**19784 The *inet\_addr()* function shall convert the string pointed to by *cp*, in the standard IPv4 dotted  
19785 decimal notation, to an integer value suitable for use as an Internet address.19786 The *inet\_ntoa()* function shall convert the Internet host address specified by *in* to a string in the  
19787 Internet standard dot notation.19788 The *inet\_ntoa()* function need not be reentrant. A function that is not required to be reentrant is  
19789 not required to be thread-safe.

19790 All Internet addresses shall be returned in network order (bytes ordered from left to right).

19791 Values specified using IPv4 dotted decimal notation take one of the following forms:

19792 a.b.c.d When four parts are specified, each shall be interpreted as a byte of data and |  
19793 assigned, from left to right, to the four bytes of an Internet address. |19794 a.b.c When a three-part address is specified, the last part shall be interpreted as a 16-bit |  
19795 quantity and placed in the rightmost two bytes of the network address. This makes |  
19796 the three-part address format convenient for specifying Class B network addresses |  
19797 as **128.net.host**. |19798 a.b When a two-part address is supplied, the last part shall be interpreted as a 24-bit |  
19799 quantity and placed in the rightmost three bytes of the network address. This |  
19800 makes the two-part address format convenient for specifying Class A network |  
19801 addresses as **net.host**. |19802 a When only one part is given, the value shall be stored directly in the network |  
19803 address without any byte rearrangement. |19804 All numbers supplied as parts in IPv4 dotted decimal notation may be decimal, octal, or  
19805 hexadecimal, as specified in the ISO C standard (that is, a leading 0x or 0X implies hexadecimal;  
19806 otherwise, a leading '0' implies octal; otherwise, the number is interpreted as decimal).19807 **RETURN VALUE**19808 Upon successful completion, *inet\_addr()* shall return the Internet address. Otherwise, it shall  
19809 return (**in\_addr\_t**)(-1).19810 The *inet\_ntoa()* function shall return a pointer to the network address in Internet standard dot  
19811 notation.19812 **ERRORS**

19813 No errors are defined.

19814 **EXAMPLES**

19815           None.

19816 **APPLICATION USAGE**

19817           The return value of *inet\_ntoa()* may point to static data that may be overwritten by subsequent  
19818           calls to *inet\_ntoa()*.

19819 **RATIONALE**

19820           None.

19821 **FUTURE DIRECTIONS**

19822           None.

19823 **SEE ALSO**19824           *endhostent()*, *endnetent()*, the Base Definitions volume of IEEE Std 1003.1-200x, <arpa/inet.h>19825 **CHANGE HISTORY**

19826           First released in Issue 6. Derived from the XNS, Issue 5.2 specification.

19827 **NAME**

19828       inet\_ntoa — IPv4 address manipulation

19829 **SYNOPSIS**

19830       #include <arpa/inet.h>

19831       char \*inet\_ntoa(struct in\_addr *in*);

19832 **DESCRIPTION**

19833       Refer to *inet\_addr()*.

## 19834 NAME

19835 inet\_ntop, inet\_pton — convert IPv4 and IPv6 addresses between binary and text form

## 19836 SYNOPSIS

19837 #include <arpa/inet.h>

19838 const char \*inet\_ntop(int af, const void \*restrict src,  
19839 char \*restrict dst, socklen\_t size);

19840 int inet\_pton(int af, const char \*restrict src, void \*restrict dst);

## 19841 DESCRIPTION

19842 The *inet\_ntop()* function shall convert a numeric address into a text string suitable for |  
19843 IP6 presentation. The *af* argument shall specify the family of the address. This can be AF\_INET or |  
19844 AF\_INET6. The *src* argument points to a buffer holding an IPv4 address if the *af* argument is  
19845 IP6 AF\_INET, or an IPv6 address if the *af* argument is AF\_INET6. The *dst* argument points to a  
19846 buffer where the function stores the resulting text string; it shall not be NULL. The *size* argument  
19847 specifies the size of this buffer, which shall be large enough to hold the text string  
19848 IP6 (INET\_ADDRSTRLEN characters for IPv4, INET6\_ADDRSTRLEN characters for IPv6).

19849 The *inet\_pton()* function shall convert an address in its standard text presentation form into its |  
19850 IP6 numeric binary form. The *af* argument shall specify the family of the address. The AF\_INET and |  
19851 AF\_INET6 address families shall be supported. The *src* argument points to the string being |  
19852 passed in. The *dst* argument points to a buffer into which the function stores the numeric  
19853 IP6 address; this shall be large enough to hold the numeric address (32 bits for AF\_INET, 128 bits for  
19854 AF\_INET6).

19855 If the *af* argument of *inet\_pton()* is AF\_INET, the *src* string shall be in the standard IPv4 dotted-  
19856 decimal form:

19857 ddd.ddd.ddd.ddd

19858 where "ddd" is a one to three digit decimal number between 0 and 255 (see *inet\_addr()*). The  
19859 *inet\_pton()* function does not accept other formats (such as the octal numbers, hexadecimal  
19860 numbers, and fewer than four numbers that *inet\_addr()* accepts).

19861 IP6 If the *af* argument of *inet\_pton()* is AF\_INET6, the *src* string shall be in one of the following  
19862 standard IPv6 text forms:

19863 1. The preferred form is "x:x:x:x:x:x:x:x", where the 'x's are the hexadecimal values  
19864 of the eight 16-bit pieces of the address. Leading zeros in individual fields can be omitted,  
19865 but there shall be at least one numeral in every field.

19866 2. A string of contiguous zero fields in the preferred form can be shown as "::". The "::"  
19867 can only appear once in an address. Unspecified addresses ("0:0:0:0:0:0:0:0") may  
19868 be represented simply as "::".

19869 3. A third form that is sometimes more convenient when dealing with a mixed environment  
19870 of IPv4 and IPv6 nodes is "x:x:x:x:x:x.d.d.d.d", where the 'x's are the  
19871 hexadecimal values of the six high-order 16-bit pieces of the address, and the 'd's are the  
19872 decimal values of the four low-order 8-bit pieces of the address (standard IPv4  
19873 representation).

19874 **Note:** A more extensive description of the standard representations of IPv6 addresses can be found in  
19875 RFC 2373.

19876

19877 **RETURN VALUE**

19878 The *inet\_ntop()* function shall return a pointer to the buffer containing the text string if the  
19879 conversion succeeds, and NULL otherwise, and set *errno* to indicate the error.

19880 The *inet\_pton()* function shall return 1 if the conversion succeeds, with the address pointed to by  
19881 *dst* in network byte order. It shall return 0 if the input is not a valid IPv4 dotted-decimal string or  
19882 a valid IPv6 address string, or -1 with *errno* set to [EAFNOSUPPORT] if the *af* argument is  
19883 unknown.

19884 **ERRORS**

19885 The *inet\_ntop()* and *inet\_pton()* functions shall fail if:

19886 [EAFNOSUPPORT]

19887 The *af* argument is invalid.

19888 [ENOSPC] The size of the *inet\_ntop()* result buffer is inadequate.

19889 **EXAMPLES**

19890 None.

19891 **APPLICATION USAGE**

19892 None.

19893 **RATIONALE**

19894 None.

19895 **FUTURE DIRECTIONS**

19896 None.

19897 **SEE ALSO**

19898 The Base Definitions volume of IEEE Std 1003.1-200x, <[arpa/inet.h](#)>

19899 **CHANGE HISTORY**

19900 First released in Issue 6. Derived from the XNS, Issue 5.2 specification.

19901 IPv6 extensions are marked.

19902 The **restrict** keyword is added to the *inet\_ntop()* and *inet\_pton()* prototypes for alignment with  
19903 the ISO/IEC 9899:1999 standard.



19904 **NAME**

19905           initstate, random, setstate, srandom — pseudo-random number functions

19906 **SYNOPSIS**

```
19907 xSI      #include <stdlib.h>
19908          char *initstate(unsigned seed, char *state, size_t size);
19909          long random(void);
19910          char *setstate(const char *state);
19911          void srandom(unsigned seed);
```

19912

19913 **DESCRIPTION**

19914       The *random()* function shall use a non-linear additive feedback random-number generator |  
 19915       employing a default state array size of 31 **long** integers to return successive pseudo-random |  
 19916       numbers in the range from 0 to  $2^{31}-1$ . The period of this random-number generator is |  
 19917       approximately  $16 \times (2^{31}-1)$ . The size of the state array determines the period of the random- |  
 19918       number generator. Increasing the state array size shall increase the period. |

19919       With 256 bytes of state information, the period of the random-number generator shall be greater |  
 19920       than  $2^{69}$ . |

19921       Like *rand()*, *random()* shall produce by default a sequence of numbers that can be duplicated by |  
 19922       calling *srandom()* with 1 as the seed. |

19923       The *srandom()* function shall initialize the current state array using the value of *seed*. |

19924       The *initstate()* and *setstate()* functions handle restarting and changing random-number |  
 19925       generators. The *initstate()* function allows a state array, pointed to by the *state* argument, to be |  
 19926       initialized for future use. The *size* argument, which specifies the size in bytes of the state array, |  
 19927       shall be used by *initstate()* to decide what type of random-number generator to use; the larger |  
 19928       the state array, the more random the numbers. Values for the amount of state information are 8, |  
 19929       32, 64, 128, and 256 bytes. Other values greater than 8 bytes are rounded down to the nearest one |  
 19930       of these values. If *initstate()* is called with  $8 \leq \textit{size} < 32$ , then *random()* shall use a simple linear |  
 19931       congruential random number generator. The *seed* argument specifies a starting point for the |  
 19932       random-number sequence and provides for restarting at the same point. The *initstate()* function |  
 19933       shall return a pointer to the previous state information array. |

19934       If *initstate()* has not been called, then *random()* shall behave as though *initstate()* had been called |  
 19935       with *seed*=1 and *size*=128. |

19936       Once a state has been initialized, *setstate()* allows switching between state arrays. The array |  
 19937       defined by the *state* argument shall be used for further random-number generation until |  
 19938       *initstate()* is called or *setstate()* is called again. The *setstate()* function shall return a pointer to the |  
 19939       previous state array. |

19940 **RETURN VALUE**

19941       If *initstate()* is called with *size* less than 8, it shall return NULL.

19942       The *random()* function shall return the generated pseudo-random number.

19943       The *srandom()* function shall not return a value.

19944       Upon successful completion, *initstate()* and *setstate()* shall return a pointer to the previous state |  
 19945       array; otherwise, a null pointer shall be returned. |

19946 **ERRORS**

19947 No errors are defined.

19948 **EXAMPLES**

19949 None.

19950 **APPLICATION USAGE**

19951 After initialization, a state array can be restarted at a different point in one of two ways:

- 19952 1. The *initstate()* function can be used, with the desired seed, state array, and size of the  
19953 array.
- 19954 2. The *setstate()* function, with the desired state, can be used, followed by *srandom()* with the  
19955 desired seed. The advantage of using both of these functions is that the size of the state  
19956 array does not have to be saved once it is initialized.

19957 Although some implementations of *random()* have written messages to standard error, such  
19958 implementations do not conform to this volume of IEEE Std 1003.1-200x.

19959 Issue 5 restores the historical behavior of this function.

19960 Threaded applications should use *rand\_r()*, *erand48()*, *nrand48()*, or *jrand48()* instead of  
19961 *random()* when an independent random number sequence in multiple threads is required.19962 **RATIONALE**

19963 None.

19964 **FUTURE DIRECTIONS**

19965 None.

19966 **SEE ALSO**19967 *drand48()*, *rand()*, the Base Definitions volume of IEEE Std 1003.1-200x, <stdlib.h>19968 **CHANGE HISTORY**

19969 First released in Issue 4, Version 2.

19970 **Issue 5**

19971 Moved from X/OPEN UNIX extension to BASE.

19972 In the DESCRIPTION, the phrase “values smaller than 8” is replaced with “values greater than  
19973 or equal to 8, or less than 32”, “*size*<8” is replaced with “ $8 \leq \textit{size} < 32$ ”, and a new first paragraph  
19974 is added to the RETURN VALUE section. A note is added to the APPLICATION USAGE  
19975 indicating that these changes restore the historical behavior of the function.

19976 **Issue 6**

19977 In the DESCRIPTION, duplicate text “For values greater than or equal to 8 ...” is removed.

19978 **NAME**

19979           insque, remque — insert or remove an element in a queue

19980 **SYNOPSIS**

19981 xSI       #include &lt;search.h&gt;

19982           void insque(void \*element, void \*pred);

19983           void remque(void \*element);

19984

19985 **DESCRIPTION**

19986       The *insque()* and *remque()* functions shall manipulate queues built from doubly-linked lists. The  
 19987       queue can be either circular or linear. An application using *insque()* or *remque()* shall ensure it  
 19988       defines a structure in which the first two members of the structure are pointers to the same type  
 19989       of structure, and any further members are application-specific. The first member of the structure  
 19990       is a forward pointer to the next entry in the queue. The second member is a backward pointer to  
 19991       the previous entry in the queue. If the queue is linear, the queue is terminated with null  
 19992       pointers. The names of the structure and of the pointer members are not subject to any special  
 19993       restriction.

19994       The *insque()* function shall insert the element pointed to by *element* into a queue immediately  
 19995       after the element pointed to by *pred*.

19996       The *remque()* function shall remove the element pointed to by *element* from a queue.

19997       If the queue is to be used as a linear list, invoking *insque(&element, NULL)*, where *element* is the  
 19998       initial element of the queue, shall initialize the forward and backward pointers of *element* to null  
 19999       pointers.

20000       If the queue is to be used as a circular list, the application shall ensure it initializes the forward  
 20001       pointer and the backward pointer of the initial element of the queue to the element's own  
 20002       address.

20003 **RETURN VALUE**20004       The *insque()* and *remque()* functions do not return a value.20005 **ERRORS**

20006       No errors are defined.

20007 **EXAMPLES**20008       **Creating a Linear Linked List**

20009       The following example creates a linear linked list.

20010       #include &lt;search.h&gt;

20011       ...

20012       struct myque element1;

20013       struct myque element2;

20014       char \*data1 = "DATA1";

20015       char \*data2 = "DATA2";

20016       ...

20017       element1.data = data1;

20018       element2.data = data2;

20019       insque (&amp;element1, NULL);

20020       insque (&amp;element2, &amp;element1);

20021 **Creating a Circular Linked List**

20022 The following example creates a circular linked list.

```
20023 #include <search.h>
20024 ...
20025 struct myque element1;
20026 struct myque element2;

20027 char *data1 = "DATA1";
20028 char *data2 = "DATA2";
20029 ...
20030 element1.data = data1;
20031 element2.data = data2;

20032 element1.fwd = &element1;
20033 element1.bck = &element1;

20034 insque (&element2, &element1);
```

20035 **Removing an Element**

20036 The following example removes the element pointed to by *element1*.

```
20037 #include <search.h>
20038 ...
20039 struct myque element1;
20040 ...
20041 remque (&element1);
```

20042 **APPLICATION USAGE**

20043 The historical implementations of these functions described the arguments as being of type  
20044 **struct qelem \*** rather than as being of type **void \*** as defined here. In those implementations,  
20045 **struct qelem** was commonly defined in **<search.h>** as:

```
20046 struct qelem {
20047     struct qelem *q_forw;
20048     struct qelem *q_back;
20049 };
```

20050 Applications using these functions, however, were never able to use this structure directly since  
20051 it provided no room for the actual data contained in the elements. Most applications defined  
20052 structures that contained the two pointers as the initial elements and also provided space for, or  
20053 pointers to, the object's data. Applications that used these functions to update more than one  
20054 type of table also had the problem of specifying two or more different structures with the same  
20055 name, if they literally used **struct qelem** as specified.

20056 As described here, the implementations were actually expecting a structure type where the first  
20057 two members were forward and backward pointers to structures. With C compilers that didn't  
20058 provide function prototypes, applications used structures as specified in the DESCRIPTION  
20059 above and the compiler did what the application expected.

20060 If this method had been carried forward with an ISO C standard compiler and the historical  
20061 function prototype, most applications would have to be modified to cast pointers to the  
20062 structures actually used to be pointers to **struct qelem** to avoid compilation warnings. By  
20063 specifying **void \*** as the argument type, applications do not need to change (unless they  
20064 specifically referenced **struct qelem** and depended on it being defined in **<search.h>**).

20065 **RATIONALE**

20066           None.

20067 **FUTURE DIRECTIONS**

20068           None.

20069 **SEE ALSO**

20070           The Base Definitions volume of IEEE Std 1003.1-200x, <**search.h**>

20071 **CHANGE HISTORY**

20072           First released in Issue 4, Version 2.

20073 **Issue 5**

20074           Moved from X/OPEN UNIX extension to BASE.

20075 **Issue 6**

20076           The DESCRIPTION is updated to avoid use of the term “must” for application requirements.

## 20077 NAME

20078        ioctl — control a STREAMS device (**STREAMS**)

## 20079 SYNOPSIS

20080 XSR        #include &lt;stropts.h&gt;

20081        int ioctl(int *fildev*, int *request*, ... /\* *arg* \*/);

20082

## 20083 DESCRIPTION

20084        The *ioctl()* function shall perform a variety of control functions on STREAMS devices. For non- |  
 20085        STREAMS devices, the functions performed by this call are unspecified. The *request* argument |  
 20086        and an optional third argument (with varying type) shall be passed to and interpreted by the  
 20087        appropriate part of the STREAM associated with *fildev*.

20088        The *fildev* argument is an open file descriptor that refers to a device.

20089        The *request* argument selects the control function to be performed and shall depend on the  
 20090        STREAMS device being addressed.

20091        The *arg* argument represents additional information that is needed by this specific STREAMS  
 20092        device to perform the requested function. The type of *arg* depends upon the particular control  
 20093        request, but it shall be either an integer or a pointer to a device-specific data structure.

20094        The *ioctl()* commands applicable to STREAMS, their arguments, and error conditions that apply  
 20095        to each individual command are described below.

20096        The following *ioctl()* commands, with error values indicated, are applicable to all STREAMS  
 20097        files:

20098        I\_PUSH        Pushes the module whose name is pointed to by *arg* onto the top of the  
 20099        current STREAM, just below the STREAM head. It then calls the *open()*  
 20100        function of the newly-pushed module.

20101        The *ioctl()* function with the I\_PUSH command shall fail if:

20102                [EINVAL]        Invalid module name.

20103                [ENXIO]        Open function of new module failed.

20104                [ENXIO]        Hangup received on *fildev*.

20105        I\_POP        Removes the module just below the STREAM head of the STREAM pointed to  
 20106        by *fildev*. The *arg* argument should be 0 in an I\_POP request.

20107        The *ioctl()* function with the I\_POP command shall fail if:

20108                [EINVAL]        No module present in the STREAM.

20109                [ENXIO]        Hangup received on *fildev*.

20110        I\_LOOK       Retrieves the name of the module just below the STREAM head of the  
 20111        STREAM pointed to by *fildev*, and places it in a character string pointed to by  
 20112        *arg*. The buffer pointed to by *arg* should be at least FMNAMESZ+1 bytes long,  
 20113        where FMNAMESZ is defined in <stropts.h>.

20114        The *ioctl()* function with the I\_LOOK command shall fail if:

20115                [EINVAL]        No module present in the STREAM.

20116        I\_FLUSH       Flushes read and/or write queues, depending on the value of *arg*. Valid *arg* |  
 20117        values are:

20118	FLUSHR	Flush all read queues.
20119	FLUSHW	Flush all write queues.
20120	FLUSHRW	Flush all read and all write queues.
20121	The <i>ioctl()</i> function with the <code>L_FLUSH</code> command shall fail if:	
20122	[EINVAL]	Invalid <i>arg</i> value.
20123	[EAGAIN] or [ENOSR]	Unable to allocate buffers for flush message.
20124		
20125	[ENXIO]	Hangup received on <i>fildev</i> .
20126	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.
20127		
20128		
20129		
20130	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:
20131		
20132		
20133		
20134		
20135		
20136	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.
20137		
20138		
20139	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.
20140		
20141		
20142	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.
20143		
20144		
20145	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.
20146		
20147		
20148	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.
20149		
20150		
20151		
20152	S_WRNORM	Equivalent to S_OUTPUT.
20153	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.
20154		
20155		
20156		
20157	S_MSG	A STREAMS signal message that contains the SIGPOLL signal has reached the front of the STREAM head read queue.
20158		
20159		
20160	S_ERROR	Notification of an error condition has reached the STREAM head.
20161		

20162		S_HANGUP	Notification of a hangup has reached the STREAM head.
20163		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.
20164			
20165			
20166			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> .
20167			
20168			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.
20169			
20170			
20171			
20172			The <i>ioctl()</i> function with the I_SETSIG command shall fail if:
20173		[EINVAL]	The value of <i>arg</i> is invalid.
20174		[EINVAL]	The value of <i>arg</i> is 0 and the calling process is not registered to receive the SIGPOLL signal.
20175			
20176		[EAGAIN]	There were insufficient resources to store the signal request.
20177	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 <i>int</i> pointed to by <i>arg</i> , where the events are those specified in the description of I_SETSIG above.
20178			
20179			
20180			
20181			The <i>ioctl()</i> function with the I_GETSIG command shall fail if:
20182		[EINVAL]	Process is not registered to receive the SIGPOLL signal.
20183	I_FIND		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.
20184			
20185			
20186			The <i>ioctl()</i> function with the I_FIND command shall fail if:
20187		[EINVAL]	<i>arg</i> does not contain a valid module name.
20188	I_PEEK		Retrieves 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 <i>strpeek</i> structure.
20189			
20190			
20191			
20192			The application shall ensure that the <i>maxlen</i> member in the <i>ctlbuf</i> and <i>databuf</i> 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.
20193			
20194			
20195			
20196			
20197			
20198			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, <i>ctlbuf</i> specifies information in the control buffer, <i>databuf</i> specifies information in the data buffer, and <i>flags</i> contains the value RS_HIPRI or 0.
20199			
20200			
20201			
20202			
20203			
20204	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:
20205			



20206		RNORM	Byte-stream mode, the default.
20207		RMSGD	Message-discard mode.
20208		RMSGN	Message-nondiscard mode.
20209			The bitwise-inclusive OR of RMSGD and RMSGN shall return [EINVAL]. The
20210			bitwise-inclusive OR of RNORM and either RMSGD or RMSGN shall result in
20211			the other flag overriding RNORM which is the default.
20212			In addition, treatment of control messages by the STREAM head may be
20213			changed by setting any of the following flags in <i>arg</i> :
20214		RPROTNORM	Fail <i>read()</i> with [EBADMSG] if a message containing a
20215			control part is at the front of the STREAM head read queue.
20216		RPROTDAT	Deliver the control part of a message as data when a
20217			process issues a <i>read()</i> .
20218		RPROTDIS	Discard the control part of a message, delivering any data
20219			portion, when a process issues a <i>read()</i> .
20220			The <i>ioctl()</i> function with the I_SRDOPT command shall fail if:
20221		[EINVAL]	The <i>arg</i> argument is not valid.
20222	I_GRDOPT		Returns the current read mode setting as, described above, in an <b>int</b> pointed to
20223			by the argument <i>arg</i> . Read modes are described in <i>read()</i> .
20224	I_NREAD		Counts the number of data bytes in the data part of the first message on the
20225			STREAM head read queue and places this value in the <b>int</b> pointed to by <i>arg</i> .
20226			The return value for the command shall be the number of messages on the
20227			STREAM head read queue. For example, if 0 is returned in <i>arg</i> , but the <i>ioctl()</i>
20228			return value is greater than 0, this indicates that a zero-length message is next
20229			on the queue.
20230	I_FDINSERT		Creates a message from specified buffer(s), adds information about another
20231			STREAM, and sends the message downstream. The message contains a
20232			control part and an optional data part. The data and control parts to be sent
20233			are distinguished by placement in separate buffers, as described below. The
20234			<i>arg</i> argument points to a <b>strfdinsert</b> structure.
20235			The application shall ensure that the <i>len</i> member in the <b>ctlbuf strbuf</b> structure
20236			is set to the size of a <b>t_uscalar_t</b> plus the number of bytes of control
20237			information to be sent with the message. The <i>fildev</i> member specifies the file
20238			descriptor of the other STREAM, and the <i>offset</i> member, which must be
20239			suitably aligned for use as a <b>t_uscalar_t</b> , specifies the offset from the start of
20240			the control buffer where I_FDINSERT shall store a <b>t_uscalar_t</b> whose
20241			interpretation is specific to the STREAM end. The application shall ensure that
20242			the <i>len</i> member in the <b>datbuf strbuf</b> structure is set to the number of bytes of
20243			data information to be sent with the message, or to 0 if no data part is to be
20244			sent.
20245			The <i>flags</i> member specifies the type of message to be created. A normal
20246			message is created if <i>flags</i> is set to 0, and a high-priority message is created if
20247			<i>flags</i> is set to RS_HIPRI. For non-priority messages, I_FDINSERT shall block if
20248			the STREAM write queue is full due to internal flow control conditions. For
20249			priority messages, I_FDINSERT does not block on this condition. For non-
20250			priority messages, I_FDINSERT does not block when the write queue is full

20251 and O\_NONBLOCK is set. Instead, it fails and sets *errno* to [EAGAIN].

20252 I\_FDINSERT also blocks, unless prevented by lack of internal resources,  
20253 waiting for the availability of message blocks in the STREAM, regardless of  
20254 priority or whether O\_NONBLOCK has been specified. No partial message is  
20255 sent.

20256 The *ioctl()* function with the I\_FDINSERT command shall fail if:

20257 [EAGAIN] A non-priority message is specified, the O\_NONBLOCK  
20258 flag is set, and the STREAM write queue is full due to  
20259 internal flow control conditions.

20260 [EAGAIN] or [ENOSR]  
20261 Buffers cannot be allocated for the message that is to be  
20262 created.

20263 [EINVAL] One of the following:

20264 — The *fildev* member of the **strfdinsert** structure is not a  
20265 valid, open STREAM file descriptor.

20266 — The size of a **t\_uscalar\_t** plus *offset* is greater than the *len*  
20267 member for the buffer specified through **ctlbuf**.

20268 — The *offset* member does not specify a properly-aligned  
20269 location in the data buffer.

20270 — An undefined value is stored in *flags*.

20271 [ENXIO] Hangupt received on the STREAM identified by either the  
20272 *fildev* argument or the *fildev* member of the **strfdinsert**  
20273 structure.

20274 [ERANGE] The *len* member for the buffer specified through **databuf**  
20275 does not fall within the range specified by the maximum  
20276 and minimum packet sizes of the topmost STREAM module  
20277 or the *len* member for the buffer specified through **databuf**  
20278 is larger than the maximum configured size of the data part  
20279 of a message; or the *len* member for the buffer specified  
20280 through **ctlbuf** is larger than the maximum configured size  
20281 of the control part of a message.

20282 I\_STR Constructs an internal STREAMS *ioctl()* message from the data pointed to by  
20283 *arg*, and sends that message downstream.

20284 This mechanism is provided to send *ioctl()* requests to downstream modules  
20285 and drivers. It allows information to be sent with *ioctl()*, and returns to the  
20286 process any information sent upstream by the downstream recipient. I\_STR  
20287 shall block until the system responds with either a positive or negative  
20288 acknowledgement message, or until the request times out after some period of  
20289 time. If the request times out, it shall fail with *errno* set to [ETIME].

20290 At most, one I\_STR can be active on a STREAM. Further I\_STR calls shall  
20291 block until the active I\_STR completes at the STREAM head. The default  
20292 timeout interval for these requests is 15 seconds. The O\_NONBLOCK flag has  
20293 no effect on this call.

20294 To send requests downstream, the application shall ensure that *arg* points to a  
20295 **striocctl** structure.

20296		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
20297		(-1=infinite, 0=use implementation-defined timeout interval, >0=as specified)
20298		an I_STR request shall wait for acknowledgement before timing out. <i>ic_len</i> is
20299		the number of bytes in the data argument, and <i>ic_dp</i> is a pointer to the data
20300		argument. The <i>ic_len</i> member has two uses: on input, it contains the length of
20301		the data argument passed in, and on return from the command, it contains the
20302		number of bytes being returned to the process (the buffer pointed to by <i>ic_dp</i>
20303		should be large enough to contain the maximum amount of data that any
20304		module or the driver in the STREAM can return).
20305		
20306		The STREAM head shall convert the information pointed to by the <b>strioc</b>
20307		structure to an internal <i>ioctl()</i> command message and sends it downstream.
20308		The <i>ioctl()</i> function with the I_STR command shall fail if:
20309		[EAGAIN] or [ENOSR]
20310		Unable to allocate buffers for the <i>ioctl()</i> message.
20311		[EINVAL] The <i>ic_len</i> member is less than 0 or larger than the
20312		maximum configured size of the data part of a message, or
20313		<i>ic_timeout</i> is less than -1.
20314		[ENXIO] Hangup received on <i>fil</i> des.
20315		[ETIME] A downstream <i>ioctl()</i> timed out before acknowledgement
20316		was received.
20317		An I_STR can also fail while waiting for an acknowledgement if a message
20318		indicating an error or a hangup is received at the STREAM head. In addition,
20319		an error code can be returned in the positive or negative acknowledgement
20320		message, in the event the <i>ioctl()</i> command sent downstream fails. For these
20321		cases, I_STR shall fail with <i>errno</i> set to the value in the message.
20322	I_SWROPT	Sets the write mode using the value of the argument <i>arg</i> . Valid bit settings for
20323		<i>arg</i> are:
20324		SNDZERO Send a zero-length message downstream when a <i>write()</i> of
20325		0 bytes occurs. To not send a zero-length message when a
20326		<i>write()</i> of 0 bytes occurs, the application shall ensure that
20327		this bit is not set in <i>arg</i> (for example, <i>arg</i> would be set to 0).
20328		The <i>ioctl()</i> function with the I_SWROPT command shall fail if:
20329		[EINVAL] <i>arg</i> is not the above value.
20330	I_GWROPT	Returns the current write mode setting, as described above, in the <b>int</b> that is
20331		pointed to by the argument <i>arg</i> .
20332	I_SENDFD	Creates a new reference to the open file description associated with the file
20333		descriptor <i>arg</i> , and writes a message on the STREAMS-based pipe <i>fil</i> des
20334		containing this reference, together with the user ID and group ID of the calling
20335		process.
20336		The <i>ioctl()</i> function with the I_SENDFD command shall fail if:
20337		[EAGAIN] The sending STREAM is unable to allocate a message block
20338		to contain the file pointer; or the read queue of the receiving
20339		STREAM head is full and cannot accept the message sent by
20340		I_SENDFD.

20341		[EBADF]	The <i>arg</i> argument is not a valid, open file descriptor.
20342		[EINVAL]	The <i>fildev</i> argument is not connected to a STREAM pipe.
20343		[ENXIO]	Hangup received on <i>fildev</i> .
20344	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> .
20345			
20346			
20347			
20348			
20349			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.
20350			
20351			If O_NONBLOCK is not set, I_RECVFD shall block until a message is present at the STREAM head. If O_NONBLOCK is set, I_RECVFD shall fail with <i>errno</i> set to [EAGAIN] if no message is present at the STREAM head.
20352			
20353			
20354			If the message at the STREAM head is a message sent by an I_SENDFD, a new file descriptor shall be 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> .
20355			
20356			
20357			
20358			The <i>ioctl()</i> function with the I_RECVFD command shall fail if:
20359		[EAGAIN]	A message is not present at the STREAM head read queue and the O_NONBLOCK flag is set.
20360			
20361		[EBADMSG]	The message at the STREAM head read queue is not a message containing a passed file descriptor.
20362			
20363		[EMFILE]	The process has the maximum number of file descriptors currently open that it is allowed.
20364			
20365		[ENXIO]	Hangup received on <i>fildev</i> .
20366	I_LIST		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 shall be 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.
20367			
20368			
20369			
20370			
20371			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 shall contain 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> shall be 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.
20372			
20373			
20374			
20375			
20376			
20377			
20378			
20379			The <i>ioctl()</i> function with the I_LIST command shall fail if:
20380		[EINVAL]	The <i>sl_nmods</i> member is less than 1.
20381		[EAGAIN] or [ENOSR]	Unable to allocate buffers.
20382			
20383	I_ATMARK		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> argument determines
20384			

20385		how the checking is done when there may be multiple marked messages on
20386		the STREAM head read queue. It may take on the following values:
20387	ANYMARK	Check if the message is marked.
20388	LASTMARK	Check if the message is the last one marked on the queue.
20389		The bitwise-inclusive OR of the flags ANYMARK and LASTMARK is
20390		permitted.
20391		The return value shall be 1 if the mark condition is satisfied; otherwise, the
20392		value shall be 0.
20393		The <i>ioctl()</i> function with the L_ATMARK command shall fail if:
20394		[EINVAL] Invalid <i>arg</i> value.
20395	I_CKBAND	Checks if the message of a given priority band exists on the STREAM head
20396		read queue. This shall return 1 if a message of the given priority exists, 0 if no
20397		such message exists, or -1 on error. <i>arg</i> should be of type <b>int</b> .
20398		The <i>ioctl()</i> function with the I_CKBAND command shall fail if:
20399		[EINVAL] Invalid <i>arg</i> value.
20400	I_GETBAND	Returns the priority band of the first message on the STREAM head read
20401		queue in the integer referenced by <i>arg</i> .
20402		The <i>ioctl()</i> function with the I_GETBAND command shall fail if:
20403		[ENODATA] No message on the STREAM head read queue.
20404	I_CANPUT	Checks if a certain band is writable. <i>arg</i> is set to the priority band in question.
20405		The return value shall be 0 if the band is flow-controlled, 1 if the band is
20406		writable, or -1 on error.
20407		The <i>ioctl()</i> function with the I_CANPUT command shall fail if:
20408		[EINVAL] Invalid <i>arg</i> value.
20409	I_SETCLTIME	This request allows the process to set the time the STREAM head shall delay
20410		when a STREAM is closing and there is data on the write queues. Before
20411		closing each module or driver, if there is data on its write queue, the STREAM
20412		head shall delay for the specified amount of time to allow the data to drain. If,
20413		after the delay, data is still present, it shall be flushed. The <i>arg</i> argument is a
20414		pointer to an integer specifying the number of milliseconds to delay, rounded
20415		up to the nearest valid value. If I_SETCLTIME is not performed on a STREAM,
20416		an implementation-defined default timeout interval is used.
20417		The <i>ioctl()</i> function with the I_SETCLTIME command shall fail if:
20418		[EINVAL] Invalid <i>arg</i> value.
20419	I_GETCLTIME	Returns the close time delay in the integer pointed to by <i>arg</i> .

20420 **Multiplexed STREAMS Configurations**

20421 The following commands are used for connecting and disconnecting multiplexed STREAMS  
20422 configurations. These commands use an implementation-defined default timeout interval.

20423 **I\_LINK** Connects two STREAMs, where *filde*s is the file descriptor of the STREAM  
20424 connected to the multiplexing driver, and *arg* is the file descriptor of the  
20425 STREAM connected to another driver. The STREAM designated by *arg* is  
20426 connected below the multiplexing driver. I\_LINK requires the multiplexing  
20427 driver to send an acknowledgement message to the STREAM head regarding  
20428 the connection. This call shall return a multiplexer ID number (an identifier  
20429 used to disconnect the multiplexer; see I\_UNLINK) on success, and -1 on  
20430 failure.

20431 The *ioctl()* function with the I\_LINK command shall fail if:

20432 [ENXIO] Hangup received on *filde*s.  
20433 [ETIME] Timeout before acknowledgement message was received at  
20434 STREAM head.  
20435 [EAGAIN] or [ENOSR]  
20436 Unable to allocate STREAMS storage to perform the  
20437 I\_LINK.  
20438 [EBADF] The *arg* argument is not a valid, open file descriptor.  
20439 [EINVAL] The *filde*s argument does not support multiplexing; or *arg* is  
20440 not a STREAM or is already connected downstream from a  
20441 multiplexer; or the specified I\_LINK operation would  
20442 connect the STREAM head in more than one place in the  
20443 multiplexed STREAM.

20444 An I\_LINK can also fail while waiting for the multiplexing driver to  
20445 acknowledge the request, if a message indicating an error or a hangup is  
20446 received at the STREAM head of *filde*s. In addition, an error code can be  
20447 returned in the positive or negative acknowledgement message. For these  
20448 cases, I\_LINK fails with *errno* set to the value in the message.

20449 **I\_UNLINK** Disconnects the two STREAMs specified by *filde*s and *arg*. *filde*s is the file  
20450 descriptor of the STREAM connected to the multiplexing driver. The *arg*  
20451 argument is the multiplexer ID number that was returned by the I\_LINK  
20452 *ioctl()* command when a STREAM was connected downstream from the  
20453 multiplexing driver. If *arg* is MUXID\_ALL, then all STREAMs that were  
20454 connected to *filde*s shall be disconnected. As in I\_LINK, this command  
20455 requires acknowledgement.

20456 The *ioctl()* function with the I\_UNLINK command shall fail if:

20457 [ENXIO] Hangup received on *filde*s.  
20458 [ETIME] Timeout before acknowledgement message was received at  
20459 STREAM head.  
20460 [EAGAIN] or [ENOSR]  
20461 Unable to allocate buffers for the acknowledgement  
20462 message.  
20463 [EINVAL] Invalid multiplexer ID number.

20464 An I\_UNLINK can also fail while waiting for the multiplexing driver to  
 20465 acknowledge the request if a message indicating an error or a hangup is  
 20466 received at the STREAM head of *filde*s. In addition, an error code can be  
 20467 returned in the positive or negative acknowledgement message. For these  
 20468 cases, I\_UNLINK shall fail with *errno* set to the value in the message.

20469 I\_PLINK Creates a *persistent connection* between two STREAMs, where *filde*s is the file  
 20470 descriptor of the STREAM connected to the multiplexing driver, and *arg* is the  
 20471 file descriptor of the STREAM connected to another driver. This call shall  
 20472 create a persistent connection which can exist even if the file descriptor *filde*s  
 20473 associated with the upper STREAM to the multiplexing driver is closed. The  
 20474 STREAM designated by *arg* gets connected via a persistent connection below  
 20475 the multiplexing driver. I\_PLINK requires the multiplexing driver to send an  
 20476 acknowledgement message to the STREAM head. This call shall return a  
 20477 multiplexer ID number (an identifier that may be used to disconnect the  
 20478 multiplexer; see I\_PUNLINK) on success, and -1 on failure.

20479 The *ioctl*() function with the I\_PLINK command shall fail if:

20480 [ENXIO] Hangup received on *filde*s.

20481 [ETIME] Timeout before acknowledgement message was received at  
 20482 STREAM head.

20483 [EAGAIN] or [ENOSR]  
 20484 Unable to allocate STREAMS storage to perform the  
 20485 I\_PLINK.

20486 [EBADF] The *arg* argument is not a valid, open file descriptor.

20487 [EINVAL] The *filde*s argument does not support multiplexing; or *arg* is  
 20488 not a STREAM or is already connected downstream from a  
 20489 multiplexer; or the specified I\_PLINK operation would  
 20490 connect the STREAM head in more than one place in the  
 20491 multiplexed STREAM.

20492 An I\_PLINK can also fail while waiting for the multiplexing driver to  
 20493 acknowledge the request, if a message indicating an error or a hangup is  
 20494 received at the STREAM head of *filde*s. In addition, an error code can be  
 20495 returned in the positive or negative acknowledgement message. For these  
 20496 cases, I\_PLINK shall fail with *errno* set to the value in the message.

20497 I\_PUNLINK Disconnects the two STREAMs specified by *filde*s and *arg* from a persistent  
 20498 connection. The *filde*s argument is the file descriptor of the STREAM  
 20499 connected to the multiplexing driver. The *arg* argument is the multiplexer ID  
 20500 number that was returned by the I\_PLINK *ioctl*() command when a STREAM  
 20501 was connected downstream from the multiplexing driver. If *arg* is  
 20502 MUXID\_ALL, then all STREAMs which are persistent connections to *filde*s  
 20503 shall be disconnected. As in I\_PLINK, this command requires the multiplexing  
 20504 driver to acknowledge the request.

20505 The *ioctl*() function with the I\_PUNLINK command shall fail if:

20506 [ENXIO] Hangup received on *filde*s.

20507 [ETIME] Timeout before acknowledgement message was received at  
 20508 STREAM head.

20509 [EAGAIN] or [ENOSR]  
 20510 Unable to allocate buffers for the acknowledgement  
 20511 message.  
 20512 [EINVAL] Invalid multiplexer ID number.  
 20513 An I\_PUNLINK can also fail while waiting for the multiplexing driver to  
 20514 acknowledge the request if a message indicating an error or a hangup is  
 20515 received at the STREAM head of *fildev*. In addition, an error code can be  
 20516 returned in the positive or negative acknowledgement message. For these  
 20517 cases, I\_PUNLINK shall fail with *errno* set to the value in the message.

20518 **RETURN VALUE**

20519 Upon successful completion, *ioctl()* shall return a value other than  $-1$  that depends upon the  
 20520 STREAMS device control function. Otherwise, it shall return  $-1$  and set *errno* to indicate the  
 20521 error.

20522 **ERRORS**

20523 Under the following general conditions, *ioctl()* shall fail if:

20524 [EBADF] The *fildev* argument is not a valid open file descriptor.  
 20525 [EINTR] A signal was caught during the *ioctl()* operation.  
 20526 [EINVAL] The STREAM or multiplexer referenced by *fildev* is linked (directly or  
 20527 indirectly) downstream from a multiplexer.

20528 If an underlying device driver detects an error, then *ioctl()* shall fail if:

20529 [EINVAL] The *request* or *arg* argument is not valid for this device.  
 20530 [EIO] Some physical I/O error has occurred.  
 20531 [ENOTTY] The *fildev* argument is not associated with a STREAMS device that accepts  
 20532 control functions.  
 20533 [ENXIO] The *request* and *arg* arguments are valid for this device driver, but the service  
 20534 requested cannot be performed on this particular sub-device.  
 20535 [ENODEV] The *fildev* argument refers to a valid STREAMS device, but the corresponding  
 20536 device driver does not support the *ioctl()* function.

20537 If a STREAM is connected downstream from a multiplexer, any *ioctl()* command except  
 20538 I\_UNLINK and I\_PUNLINK shall set *errno* to [EINVAL].

20539 **EXAMPLES**

20540 None.

20541 **APPLICATION USAGE**

20542 The implementation-defined timeout interval for STREAMS has historically been 15 seconds.

20543 **RATIONALE**

20544 None.

20545 **FUTURE DIRECTIONS**

20546 None.

20547 **SEE ALSO**

20548 *close()*, *fcntl()*, *getmsg()*, *open()*, *pipe()*, *poll()*, *putmsg()*, *read()*, *sigaction()*, *write()*, the Base  
 20549 Definitions volume of IEEE Std 1003.1-200x, <**stropts.h**>, Section 2.6 (on page 488)



20550 **CHANGE HISTORY**

20551 First released in Issue 4, Version 2.

20552 **Issue 5**

20553 Moved from X/OPEN UNIX extension to BASE.

20554 **Issue 6**

20555 The Open Group Corrigendum U028/4 is applied, correcting text in the I\_FDINSERT, [EINVAL]  
20556 case to refer to *ctlbuf*.

20557 This function is marked as part of the XSI STREAMS Option Group.

20558 The DESCRIPTION is updated to avoid use of the term “must” for application requirements.

20559 **NAME**

20560 isalnum — test for an alphanumeric character

20561 **SYNOPSIS**

20562 #include <ctype.h>

20563 int isalnum(int c);

20564 **DESCRIPTION**

20565 cx The functionality described on this reference page is aligned with the ISO C standard. Any  
20566 conflict between the requirements described here and the ISO C standard is unintentional. This  
20567 volume of IEEE Std 1003.1-200x defers to the ISO C standard.

20568 The *isalnum()* function shall test whether *c* is a character of class **alpha** or **digit** in the program's  
20569 current locale; see the Base Definitions volume of IEEE Std 1003.1-200x, Chapter 7, Locale.

20570 The *c* argument is an **int**, the value of which the application shall ensure is representable as an  
20571 **unsigned char** or equal to the value of the macro EOF. If the argument has any other value, the  
20572 behavior is undefined.

20573 **RETURN VALUE**

20574 The *isalnum()* function shall return non-zero if *c* is an alphanumeric character; otherwise, it shall  
20575 return 0.

20576 **ERRORS**

20577 No errors are defined.

20578 **EXAMPLES**

20579 None.

20580 **APPLICATION USAGE**

20581 To ensure applications portability, especially across natural languages, only this function and  
20582 those listed in the SEE ALSO section should be used for character classification.

20583 **RATIONALE**

20584 None.

20585 **FUTURE DIRECTIONS**

20586 None.

20587 **SEE ALSO**

20588 *isalpha()*, *iscntrl()*, *isdigit()*, *isgraph()*, *islower()*, *isprint()*, *ispunct()*, *isspace()*, *isupper()*, *isxdigit()*,  
20589 *setlocale()*, the Base Definitions volume of IEEE Std 1003.1-200x, <ctype.h>, <stdio.h>, the Base  
20590 Definitions volume of IEEE Std 1003.1-200x, Chapter 7, Locale

20591 **CHANGE HISTORY**

20592 First released in Issue 1. Derived from Issue 1 of the SVID.

20593 **Issue 6**

20594 The DESCRIPTION is updated to avoid use of the term “must” for application requirements.

20595 **NAME**

20596           isalpha — test for an alphabetic character

20597 **SYNOPSIS**

20598           #include <ctype.h>

20599           int isalpha(int c);

20600 **DESCRIPTION**

20601 cx       The functionality described on this reference page is aligned with the ISO C standard. Any  
20602       conflict between the requirements described here and the ISO C standard is unintentional. This  
20603       volume of IEEE Std 1003.1-200x defers to the ISO C standard.

20604       The *isalpha()* function shall test whether *c* is a character of class **alpha** in the program's current  
20605       locale; see the Base Definitions volume of IEEE Std 1003.1-200x, Chapter 7, Locale.

20606       The *c* argument is an **int**, the value of which the application shall ensure is representable as an  
20607       **unsigned char** or equal to the value of the macro EOF. If the argument has any other value, the  
20608       behavior is undefined.

20609 **RETURN VALUE**

20610       The *isalpha()* function shall return non-zero if *c* is an alphabetic character; otherwise, it shall  
20611       return 0.

20612 **ERRORS**

20613       No errors are defined.

20614 **EXAMPLES**

20615       None.

20616 **APPLICATION USAGE**

20617       To ensure applications portability, especially across natural languages, only this function and  
20618       those listed in the SEE ALSO section should be used for character classification.

20619 **RATIONALE**

20620       None.

20621 **FUTURE DIRECTIONS**

20622       None.

20623 **SEE ALSO**

20624       *isalnum()*, *isctrl()*, *isdigit()*, *isgraph()*, *islower()*, *isprint()*, *ispunct()*, *isspace()*, *isupper()*,  
20625       *isxdigit()*, *setlocale()*, the Base Definitions volume of IEEE Std 1003.1-200x, <ctype.h>, <stdio.h>,  
20626       the Base Definitions volume of IEEE Std 1003.1-200x, Chapter 7, Locale

20627 **CHANGE HISTORY**

20628       First released in Issue 1. Derived from Issue 1 of the SVID.

20629 **Issue 6**

20630       The DESCRIPTION is updated to avoid use of the term “must” for application requirements.

20631 **NAME**20632            **isascii** — test for a 7-bit US-ASCII character20633 **SYNOPSIS**

20634 xSI        #include &lt;ctype.h&gt;

20635            int isascii(int c);

20636

20637 **DESCRIPTION**20638            The *isascii()* function shall test whether *c* is a 7-bit US-ASCII character code.20639            The *isascii()* function is defined on all integer values.20640 **RETURN VALUE**20641            The *isascii()* function shall return non-zero if *c* is a 7-bit US-ASCII character code between 0 and

20642            octal 0177 inclusive; otherwise, it shall return 0.

20643 **ERRORS**

20644            No errors are defined.

20645 **EXAMPLES**

20646            None.

20647 **APPLICATION USAGE**

20648            None.

20649 **RATIONALE**

20650            None.

20651 **FUTURE DIRECTIONS**

20652            None.

20653 **SEE ALSO**

20654            The Base Definitions volume of IEEE Std 1003.1-200x, &lt;ctype.h&gt;

20655 **CHANGE HISTORY**

20656            First released in Issue 1. Derived from Issue 1 of the SVID.

20657 **NAME**20658           isastream — test a file descriptor (**STREAMS**)20659 **SYNOPSIS**20660 *xsr*       #include <stropts.h>20661           int isastream(int *fildev*);

20662

20663 **DESCRIPTION**20664           The *isastream()* function shall test whether *fildev*, an open file descriptor, is associated with a  
20665           STREAMS-based file.20666 **RETURN VALUE**20667           Upon successful completion, *isastream()* shall return 1 if *fildev* refers to a STREAMS-based file  
20668           and 0 if not. Otherwise, *isastream()* shall return -1 and set *errno* to indicate the error.20669 **ERRORS**20670           The *isastream()* function shall fail if:20671           [EBADF]           The *fildev* argument is not a valid open file descriptor.20672 **EXAMPLES**

20673           None.

20674 **APPLICATION USAGE**

20675           None.

20676 **RATIONALE**

20677           None.

20678 **FUTURE DIRECTIONS**

20679           None.

20680 **SEE ALSO**

20681           The Base Definitions volume of IEEE Std 1003.1-200x, &lt;stropts.h&gt;

20682 **CHANGE HISTORY**

20683           First released in Issue 4, Version 2.

20684 **Issue 5**

20685           Moved from X/OPEN UNIX extension to BASE.

20686 **NAME**

20687           isatty — test for a terminal device

20688 **SYNOPSIS**

20689           #include <unistd.h>

20690           int isatty(int *fdes*);

20691 **DESCRIPTION**

20692           The *isatty()* function shall test whether *fdes*, an open file descriptor, is associated with a  
20693           terminal device.

20694 **RETURN VALUE**

20695           The *isatty()* function shall return 1 if *fdes* is associated with a terminal; otherwise, it shall return  
20696           0 and may set *errno* to indicate the error.

20697 **ERRORS**

20698           The *isatty()* function may fail if:

20699           [EBADF]           The *fdes* argument is not a valid open file descriptor.

20700           [ENOTTY]         The *fdes* argument is not associated with a terminal.

20701 **EXAMPLES**

20702           None.

20703 **APPLICATION USAGE**

20704           The *isatty()* function does not necessarily indicate that a human being is available for interaction  
20705           via *fdes*. It is quite possible that non-terminal devices are connected to the communications  
20706           line.

20707 **RATIONALE**

20708           None.

20709 **FUTURE DIRECTIONS**

20710           None.

20711 **SEE ALSO**

20712           The Base Definitions volume of IEEE Std 1003.1-200x, <unistd.h>

20713 **CHANGE HISTORY**

20714           First released in Issue 1. Derived from Issue 1 of the SVID.

20715 **Issue 6**

20716           The following new requirements on POSIX implementations derive from alignment with the  
20717           Single UNIX Specification:

- 20718           • The optional setting of *errno* to indicate an error is added.
- 20719           • The [EBADF] and [ENOTTY] optional error conditions are added.

20720 **NAME**

20721           isblank — test for a blank character

20722 **SYNOPSIS**

20723           #include <ctype.h>

20724           int isblank(int c);

20725 **DESCRIPTION**

20726 **CX**       The functionality described on this reference page is aligned with the ISO C standard. Any  
20727       conflict between the requirements described here and the ISO C standard is unintentional. This  
20728       volume of IEEE Std 1003.1-200x defers to the ISO C standard.

20729       The *isblank()* function shall test whether *c* is a character of class **blank** in the program's current  
20730       locale; see the Base Definitions volume of IEEE Std 1003.1-200x, Chapter 7, Locale.

20731       The *c* argument is a type **int**, the value of which the application shall ensure is a character  
20732       representable as an **unsigned char** or equal to the value of the macro EOF. If the argument has  
20733       any other value, the behavior is undefined.

20734 **RETURN VALUE**

20735       The *isblank()* function shall return non-zero if *c* is a <blank>; otherwise, it shall return 0.

20736 **ERRORS**

20737       No errors are defined.

20738 **EXAMPLES**

20739       None.

20740 **APPLICATION USAGE**

20741       To ensure applications portability, especially across natural languages, only this function and  
20742       those listed in the SEE ALSO section should be used for character classification.

20743 **RATIONALE**

20744       None.

20745 **FUTURE DIRECTIONS**

20746       None.

20747 **SEE ALSO**

20748       *isalnum()*, *isalpha()*, *iscntrl()*, *isdigit()*, *isgraph()*, *islower()*, *isprint()*, *ispunct()*, *isspace()*, *isupper()*,  
20749       *isxdigit()*, *setlocale()*, the Base Definitions volume of IEEE Std 1003.1-200x, <ctype.h>

20750 **CHANGE HISTORY**

20751       First released in Issue 6. Derived from the ISO/IEC 9899:1999 standard.

20752 **NAME**

20753 isctrl — test for a control character

20754 **SYNOPSIS**

20755 #include &lt;ctype.h&gt;

20756 int isctrl(int c);

20757 **DESCRIPTION**

20758 **CX** The functionality described on this reference page is aligned with the ISO C standard. Any  
20759 conflict between the requirements described here and the ISO C standard is unintentional. This  
20760 volume of IEEE Std 1003.1-200x defers to the ISO C standard.

20761 The *isctrl()* function shall test whether *c* is a character of class **cntrl** in the program's current  
20762 locale; see the Base Definitions volume of IEEE Std 1003.1-200x, Chapter 7, Locale.

20763 The *c* argument is a type **int**, the value of which the application shall ensure is a character  
20764 representable as an **unsigned char** or equal to the value of the macro EOF. If the argument has  
20765 any other value, the behavior is undefined.

20766 **RETURN VALUE**20767 The *isctrl()* function shall return non-zero if *c* is a control character; otherwise, it shall return 0.20768 **ERRORS**

20769 No errors are defined.

20770 **EXAMPLES**

20771 None.

20772 **APPLICATION USAGE**

20773 To ensure applications portability, especially across natural languages, only this function and  
20774 those listed in the SEE ALSO section should be used for character classification.

20775 **RATIONALE**

20776 None.

20777 **FUTURE DIRECTIONS**

20778 None.

20779 **SEE ALSO**

20780 *isalnum()*, *isalpha()*, *isdigit()*, *isgraph()*, *islower()*, *isprint()*, *ispunct()*, *isspace()*, *isupper()*,  
20781 *isxdigit()*, *setlocale()*, the Base Definitions volume of IEEE Std 1003.1-200x, <ctype.h>, the Base  
20782 Definitions volume of IEEE Std 1003.1-200x, Chapter 7, Locale

20783 **CHANGE HISTORY**

20784 First released in Issue 1. Derived from Issue 1 of the SVID.

20785 **Issue 6**

20786 The DESCRIPTION is updated to avoid use of the term “must” for application requirements.



20787 **NAME**

20788            isdigit — test for a decimal digit

20789 **SYNOPSIS**

20790            #include <ctype.h>

20791            int isdigit(int c);

20792 **DESCRIPTION**

20793 **CX**        The functionality described on this reference page is aligned with the ISO C standard. Any  
20794 conflict between the requirements described here and the ISO C standard is unintentional. This  
20795 volume of IEEE Std 1003.1-200x defers to the ISO C standard.

20796            The *isdigit()* function shall test whether *c* is a character of class **digit** in the program's current  
20797 locale; see the Base Definitions volume of IEEE Std 1003.1-200x, Chapter 7, Locale.

20798            The *c* argument is an **int**, the value of which the application shall ensure is a character  
20799 representable as an **unsigned char** or equal to the value of the macro EOF. If the argument has  
20800 any other value, the behavior is undefined.

20801 **RETURN VALUE**

20802            The *isdigit()* function shall return non-zero if *c* is a decimal digit; otherwise, it shall return 0.

20803 **ERRORS**

20804            No errors are defined.

20805 **EXAMPLES**

20806            None.

20807 **APPLICATION USAGE**

20808            To ensure applications portability, especially across natural languages, only this function and  
20809 those listed in the SEE ALSO section should be used for character classification.

20810 **RATIONALE**

20811            None.

20812 **FUTURE DIRECTIONS**

20813            None.

20814 **SEE ALSO**

20815            *isalnum()*, *isalpha()*, *iscntrl()*, *isgraph()*, *islower()*, *isprint()*, *ispunct()*, *isspace()*, *isupper()*,  
20816 *isxdigit()*, the Base Definitions volume of IEEE Std 1003.1-200x, <ctype.h>

20817 **CHANGE HISTORY**

20818            First released in Issue 1. Derived from Issue 1 of the SVID.

20819 **Issue 6**

20820            The DESCRIPTION is updated to avoid use of the term “must” for application requirements.

20821 **NAME**

20822           isfinite — test for finite value

20823 **SYNOPSIS**

20824           #include &lt;math.h&gt;

20825           int isfinite(real-floating x);

20826 **DESCRIPTION**

20827 cx       The functionality described on this reference page is aligned with the ISO C standard. Any  
20828       conflict between the requirements described here and the ISO C standard is unintentional. This  
20829       volume of IEEE Std 1003.1-200x defers to the ISO C standard.

20830       The *isfinite()* macro shall determine whether its argument has a finite value (zero, subnormal, or  
20831       normal, and not infinite or NaN). First, an argument represented in a format wider than its  
20832       semantic type is converted to its semantic type. Then determination is based on the type of the  
20833       argument.

20834 **RETURN VALUE**20835       The *isfinite()* macro shall return a non-zero value if and only if its argument has a finite value.20836 **ERRORS**

20837       No errors are defined.

20838 **EXAMPLES**

20839       None.

20840 **APPLICATION USAGE**

20841       None.

20842 **RATIONALE**

20843       None.

20844 **FUTURE DIRECTIONS**

20845       None.

20846 **SEE ALSO**

20847       *fpclassify()*, *isinf()*, *isnan()*, *isnormal()*, *signbit()*, the Base Definitions volume of  
20848       IEEE Std 1003.1-200x <math.h>

20849 **CHANGE HISTORY**

20850       First released in Issue 6. Derived from the ISO/IEC 9899:1999 standard.

20851 **NAME**

20852           isgraph — test for a visible character

20853 **SYNOPSIS**

20854           #include &lt;ctype.h&gt;

20855           int isgraph(int c);

20856 **DESCRIPTION**

20857 *cx*       The functionality described on this reference page is aligned with the ISO C standard. Any  
20858       conflict between the requirements described here and the ISO C standard is unintentional. This  
20859       volume of IEEE Std 1003.1-200x defers to the ISO C standard.

20860       The *isgraph()* function shall test whether *c* is a character of class **graph** in the program's current  
20861       locale; see the Base Definitions volume of IEEE Std 1003.1-200x, Chapter 7, Locale.

20862       The *c* argument is an **int**, the value of which the application shall ensure is a character  
20863       representable as an **unsigned char** or equal to the value of the macro EOF. If the argument has  
20864       any other value, the behavior is undefined.

20865 **RETURN VALUE**

20866       The *isgraph()* function shall return non-zero if *c* is a character with a visible representation;  
20867       otherwise, it shall return 0.

20868 **ERRORS**

20869       No errors are defined.

20870 **EXAMPLES**

20871       None.

20872 **APPLICATION USAGE**

20873       To ensure applications portability, especially across natural languages, only this function and  
20874       those listed in the SEE ALSO section should be used for character classification.

20875 **RATIONALE**

20876       None.

20877 **FUTURE DIRECTIONS**

20878       None.

20879 **SEE ALSO**

20880       *isalnum()*, *isalpha()*, *iscntrl()*, *isdigit()*, *islower()*, *isprint()*, *ispunct()*, *isspace()*, *isupper()*, *isxdigit()*,  
20881       *setlocale()*, the Base Definitions volume of IEEE Std 1003.1-200x, <ctype.h>, the Base Definitions  
20882       volume of IEEE Std 1003.1-200x, Chapter 7, Locale

20883 **CHANGE HISTORY**

20884       First released in Issue 1. Derived from Issue 1 of the SVID.

20885 **Issue 6**

20886       The DESCRIPTION is updated to avoid use of the term “must” for application requirements.

20887 **NAME**20888            **isgreater** — test if *x* greater than *y*20889 **SYNOPSIS**

20890            #include &lt;math.h&gt;

20891            int isgreater(real-floating *x*, real-floating *y*);20892 **DESCRIPTION**

20893 *CX*        The functionality described on this reference page is aligned with the ISO C standard. Any  
20894 conflict between the requirements described here and the ISO C standard is unintentional. This  
20895 volume of IEEE Std 1003.1-200x defers to the ISO C standard.

20896        The *isgreater()* macro shall determine whether its first argument is greater than its second  
20897 argument. The value of *isgreater(x, y)* shall be equal to  $(x) > (y)$ ; however, unlike  $(x) > (y)$ ,  
20898 *isgreater(x, y)* shall not raise the invalid floating-point exception when *x* and *y* are unordered.

20899 **RETURN VALUE**20900            Upon successful completion, the *isgreater()* macro shall return the value of  $(x) > (y)$ .20901            If *x* or *y* is NaN, 0 shall be returned.20902 **ERRORS**

20903            No errors are defined.

20904 **EXAMPLES**

20905            None.

20906 **APPLICATION USAGE**

20907        The relational and equality operators support the usual mathematical relationships between  
20908 numeric values. For any ordered pair of numeric values, exactly one of the relationships (less,  
20909 greater, and equal) is true. Relational operators may raise the invalid floating-point exception  
20910 when argument values are NaNs. For a NaN and a numeric value, or for two NaNs, just the  
20911 unordered relationship is true. This macro is a quiet (non-floating-point exception raising)  
20912 version of a relational operator. It facilitates writing efficient code that accounts for NaNs  
20913 without suffering the invalid floating-point exception. In the SYNOPSIS section, **real-floating**  
20914 indicates that the argument shall be an expression of **real-floating** type.

20915 **RATIONALE**

20916            None.

20917 **FUTURE DIRECTIONS**

20918            None.

20919 **SEE ALSO**

20920            *isgreaterequal()*, *isless()*, *islessequal()*, *islessgreater()*, *isunordered()*, the Base Definitions volume of  
20921 IEEE Std 1003.1-200x <math.h>

20922 **CHANGE HISTORY**

20923            First released in Issue 6. Derived from the ISO/IEC 9899:1999 standard.

20924 **NAME**

20925           isgreaterequal — test if *x* greater than or equal to *y*

20926 **SYNOPSIS**

20927           #include <math.h>

20928           int isgreaterequal(real-floating *x*, real-floating *y*);

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 *isgreaterequal()* macro shall determine whether its first argument is greater than or equal to  
20934 its second argument. The value of *isgreaterequal(x, y)* shall be equal to  $(x) \geq (y)$ ; however, unlike  
20935  $(x) \geq (y)$ , *isgreaterequal(x, y)* shall not raise the invalid floating-point exception when *x* and *y* are  
20936 unordered.

20937 **RETURN VALUE**

20938       Upon successful completion, the *isgreaterequal()* macro shall return the value of  $(x) \geq (y)$ .

20939       If *x* or *y* is NaN, 0 shall be returned.

20940 **ERRORS**

20941       No errors are defined.

20942 **EXAMPLES**

20943       None.

20944 **APPLICATION USAGE**

20945       The relational and equality operators support the usual mathematical relationships between  
20946 numeric values. For any ordered pair of numeric values, exactly one of the relationships (less,  
20947 greater, and equal) is true. Relational operators may raise the invalid floating-point exception  
20948 when argument values are NaNs. For a NaN and a numeric value, or for two NaNs, just the  
20949 unordered relationship is true. This macro is a quiet (non-floating-point exception raising)  
20950 version of a relational operator. It facilitates writing efficient code that accounts for NaNs  
20951 without suffering the invalid floating-point exception. In the SYNOPSIS section, **real-floating**  
20952 indicates that the argument shall be an expression of **real-floating** type.

20953 **RATIONALE**

20954       None.

20955 **FUTURE DIRECTIONS**

20956       None.

20957 **SEE ALSO**

20958       *isgreater()*, *isless()*, *islessequal()*, *islessgreater()*, *isunordered()*, the Base Definitions volume of  
20959 IEEE Std 1003.1-200x <math.h>

20960 **CHANGE HISTORY**

20961       First released in Issue 6. Derived from the ISO/IEC 9899:1999 standard.

20962 **NAME**

20963           isinf — test for infinity

20964 **SYNOPSIS**

20965           #include &lt;math.h&gt;

20966           int isinf(real-floating x);

20967 **DESCRIPTION**

20968 cx       The functionality described on this reference page is aligned with the ISO C standard. Any  
20969       conflict between the requirements described here and the ISO C standard is unintentional. This  
20970       volume of IEEE Std 1003.1-200x defers to the ISO C standard.

20971       The *isinf()* macro shall determine whether its argument value is an infinity (positive or  
20972       negative). First, an argument represented in a format wider than its semantic type is converted  
20973       to its semantic type. Then determination is based on the type of the argument.

20974 **RETURN VALUE**20975       The *isinf()* macro shall return a non-zero value if and only if its argument has an infinite value.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       *fpclassify()*, *isfinite()*, *isnan()*, *isnormal()*, *signbit()*, the Base Definitions volume of  
20988       IEEE Std 1003.1-200x <math.h>

20989 **CHANGE HISTORY**

20990       First released in Issue 6. Derived from the ISO/IEC 9899:1999 standard.

20991 **NAME**20992           isless — test if *x* is less than *y*20993 **SYNOPSIS**

20994           #include &lt;math.h&gt;

20995           int isless(real-floating *x*, real-floating *y*);20996 **DESCRIPTION**

20997 *cx*       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.

21000       The *isless()* macro shall determine whether its first argument is less than its second argument. The value of *isless(x, y)* shall be equal to  $(x) < (y)$ ; however, unlike  $(x) < (y)$ , *isless(x, y)* shall not raise the invalid floating-point exception when *x* and *y* are unordered.

21003 **RETURN VALUE**21004           Upon successful completion, the *isless()* macro shall return the value of  $(x) < (y)$ .21005           If *x* or *y* is NaN, 0 shall be returned.21006 **ERRORS**

21007           No errors are defined.

21008 **EXAMPLES**

21009           None.

21010 **APPLICATION USAGE**

21011       The relational and equality operators support the usual mathematical relationships between numeric values. For any ordered pair of numeric values, exactly one of the relationships (less, greater, and equal) is true. Relational operators may raise the invalid floating-point exception when argument values are NaNs. For a NaN and a numeric value, or for two NaNs, just the unordered relationship is true. This macro is a quiet (non-floating-point exception raising) version of a relational operator. It facilitates writing efficient code that accounts for NaNs without suffering the invalid floating-point exception. In the SYNOPSIS section, **real-floating** indicates that the argument shall be an expression of **real-floating** type.

21019 **RATIONALE**

21020           None.

21021 **FUTURE DIRECTIONS**

21022           None.

21023 **SEE ALSO**

21024           *isgreater()*, *isgreaterequal()*, *islessequal()*, *islessgreater()*, *isunordered()*, the Base Definitions volume of IEEE Std 1003.1-200x, <math.h>

21026 **CHANGE HISTORY**

21027           First released in Issue 6. Derived from the ISO/IEC 9899:1999 standard.

21028 **NAME**

21029 islessequal — test if *x* is less than or equal to *y*

21030 **SYNOPSIS**

21031 #include <math.h>

21032 int islessequal(real-floating *x*, real-floating *y*);

21033 **DESCRIPTION**

21034 **CX** The functionality described on this reference page is aligned with the ISO C standard. Any  
21035 conflict between the requirements described here and the ISO C standard is unintentional. This  
21036 volume of IEEE Std 1003.1-200x defers to the ISO C standard.

21037 The *islessequal()* macro shall determine whether its first argument is less than or equal to its  
21038 second argument. The value of *islessequal(x, y)* shall be equal to  $(x) \leq (y)$ ; however, unlike  
21039  $(x) \leq (y)$ , *islessequal(x, y)* shall not raise the invalid floating-point exception when *x* and *y* are  
21040 unordered.

21041 **RETURN VALUE**

21042 Upon successful completion, the *islessequal()* macro shall return the value of  $(x) \leq (y)$ .

21043 If *x* or *y* is NaN, 0 shall be returned.

21044 **ERRORS**

21045 No errors are defined.

21046 **EXAMPLES**

21047 None.

21048 **APPLICATION USAGE**

21049 The relational and equality operators support the usual mathematical relationships between  
21050 numeric values. For any ordered pair of numeric values, exactly one of the relationships (less,  
21051 greater, and equal) is true. Relational operators may raise the invalid floating-point exception  
21052 when argument values are NaNs. For a NaN and a numeric value, or for two NaNs, just the  
21053 unordered relationship is true. This macro is a quiet (non-floating-point exception raising)  
21054 version of a relational operator. It facilitates writing efficient code that accounts for NaNs  
21055 without suffering the invalid floating-point exception. In the SYNOPSIS section, **real-floating**  
21056 indicates that the argument shall be an expression of **real-floating** type.

21057 **RATIONALE**

21058 None.

21059 **FUTURE DIRECTIONS**

21060 None.

21061 **SEE ALSO**

21062 *isgreater()*, *isgreaterequal()*, *isless()*, *islessgreater()*, *isunordered()*, the Base Definitions volume of  
21063 IEEE Std 1003.1-200x <math.h>

21064 **CHANGE HISTORY**

21065 First released in Issue 6. Derived from the ISO/IEC 9899:1999 standard.



21066 **NAME**

21067 `islessgreater` — test if  $x$  is less than or greater than  $y$

21068 **SYNOPSIS**

21069 `#include <math.h>`

21070 `int islessgreater(real-floating  $x$ , real-floating  $y$ );`

21071 **DESCRIPTION**

21072 **CX** The functionality described on this reference page is aligned with the ISO C standard. Any  
21073 conflict between the requirements described here and the ISO C standard is unintentional. This  
21074 volume of IEEE Std 1003.1-200x defers to the ISO C standard.

21075 The `islessgreater()` macro shall determine whether its first argument is less than or greater than  
21076 its second argument. The `islessgreater( $x$ ,  $y$ )` macro is similar to  $(x) < (y) \mid \mid (x) > (y)$ ; however,  
21077 `islessgreater( $x$ ,  $y$ )` shall not raise the invalid floating-point exception when  $x$  and  $y$  are unordered  
21078 (nor shall it evaluate  $x$  and  $y$  twice).

21079 **RETURN VALUE**

21080 Upon successful completion, the `islessgreater()` macro shall return the value of  
21081  $(x) < (y) \mid \mid (x) > (y)$ .

21082 If  $x$  or  $y$  is NaN, 0 shall be returned.

21083 **ERRORS**

21084 No errors are defined.

21085 **EXAMPLES**

21086 None.

21087 **APPLICATION USAGE**

21088 The relational and equality operators support the usual mathematical relationships between  
21089 numeric values. For any ordered pair of numeric values, exactly one of the relationships (less,  
21090 greater, and equal) is true. Relational operators may raise the invalid floating-point exception  
21091 when argument values are NaNs. For a NaN and a numeric value, or for two NaNs, just the  
21092 unordered relationship is true. This macro is a quiet (non-floating-point exception raising)  
21093 version of a relational operator. It facilitates writing efficient code that accounts for NaNs  
21094 without suffering the invalid floating-point exception. In the SYNOPSIS section, **real-floating**  
21095 indicates that the argument shall be an expression of **real-floating** type.

21096 **RATIONALE**

21097 None.

21098 **FUTURE DIRECTIONS**

21099 None.

21100 **SEE ALSO**

21101 `isgreater()`, `isgreaterequal()`, `isless()`, `islessequal()`, `isunordered()`, the Base Definitions volume of  
21102 IEEE Std 1003.1-200x **<math.h>**

21103 **CHANGE HISTORY**

21104 First released in Issue 6. Derived from the ISO/IEC 9899:1999 standard.

21105 **NAME**

21106           islower — test for a lowercase letter

21107 **SYNOPSIS**

21108           #include &lt;ctype.h&gt;

21109           int islower(int c);

21110 **DESCRIPTION**

21111 cx       The functionality described on this reference page is aligned with the ISO C standard. Any  
21112       conflict between the requirements described here and the ISO C standard is unintentional. This  
21113       volume of IEEE Std 1003.1-200x defers to the ISO C standard.

21114       The *islower()* function shall test whether *c* is a character of class **lower** in the program's current  
21115       locale; see the Base Definitions volume of IEEE Std 1003.1-200x, Chapter 7, Locale.

21116       The *c* argument is an **int**, the value of which the application shall ensure is a character  
21117       representable as an **unsigned char** or equal to the value of the macro EOF. If the argument has  
21118       any other value, the behavior is undefined.

21119 **RETURN VALUE**21120       The *islower()* function shall return non-zero if *c* is a lowercase letter; otherwise, it shall return 0.21121 **ERRORS**

21122       No errors are defined.

21123 **EXAMPLES**21124       **Testing for a Lowercase Letter**

21125       The following example tests whether the value is a lowercase letter, based on the locale of the  
21126       user, then uses it as part of a key value.

```
21127       #include <ctype.h>
21128       #include <stdlib.h>
21129       #include <locale.h>
21130       ...
21131       char *keyst;
21132       int elementlen, len;
21133       char c;
21134       ...
21135       setlocale(LC_ALL, "");
21136       ...
21137       len = 0;
21138       while (len < elementlen) {
21139           c = (char) (rand() % 256);
21140       ...
21141           if (islower(c))
21142               keyst[len++] = c;
21143       }
21144       ...
```

21145 **APPLICATION USAGE**

21146       To ensure applications portability, especially across natural languages, only this function and  
21147       those listed in the SEE ALSO section should be used for character classification.

21148 **RATIONALE**

21149 None.

21150 **FUTURE DIRECTIONS**

21151 None.

21152 **SEE ALSO**21153 *isalnum()*, *isalpha()*, *iscntrl()*, *isdigit()*, *isgraph()*, *isprint()*, *ispunct()*, *isspace()*, *isupper()*,21154 *isxdigit()*, *setlocale()*, the Base Definitions volume of IEEE Std 1003.1-200x, <ctype.h>, the Base

21155 Definitions volume of IEEE Std 1003.1-200x, Chapter 7, Locale

21156 **CHANGE HISTORY**

21157 First released in Issue 1. Derived from Issue 1 of the SVID.

21158 **Issue 6**

21159 The DESCRIPTION is updated to avoid use of the term “must” for application requirements. |

21160 An example is added. |

21161 **NAME**

21162           isnan — test for a NaN

21163 **SYNOPSIS**

21164           #include &lt;math.h&gt;

21165           int isnan(real-floating x);

21166 **DESCRIPTION**

21167 **CX**       The functionality described on this reference page is aligned with the ISO C standard. Any  
21168       conflict between the requirements described here and the ISO C standard is unintentional. This  
21169       volume of IEEE Std 1003.1-200x defers to the ISO C standard.

21170       The *isnan()* macro shall determine whether its argument value is a NaN. First, an argument  
21171       represented in a format wider than its semantic type is converted to its semantic type. Then  
21172       determination is based on the type of the argument.

21173 **RETURN VALUE**21174       The *isnan()* macro shall return a non-zero value if and only if its argument has a NaN value.21175 **ERRORS**

21176       No errors are defined.

21177 **EXAMPLES**

21178       None.

21179 **APPLICATION USAGE**

21180       None.

21181 **RATIONALE**

21182       None.

21183 **FUTURE DIRECTIONS**

21184       None.

21185 **SEE ALSO**

21186       *fpclassify()*, *isfinite()*, *isinf()*, *isnormal()*, *signbit()*, the Base Definitions volume of  
21187       IEEE Std 1003.1-200x, <math.h>

21188 **CHANGE HISTORY**

21189       First released in Issue 3.

21190 **Issue 5**

21191       The DESCRIPTION is updated to indicate the return value when NaN is not supported. This  
21192       text was previously published in the APPLICATION USAGE section.

21193 **Issue 6**

21194       Entry re-written for alignment with the ISO/IEC 9899:1999 standard.

21195 **NAME**

21196           isnormal — test for a normal value

21197 **SYNOPSIS**

21198           #include &lt;math.h&gt;

21199           int isnormal(real-floating x);

21200 **DESCRIPTION**

21201 cx       The functionality described on this reference page is aligned with the ISO C standard. Any  
21202       conflict between the requirements described here and the ISO C standard is unintentional. This  
21203       volume of IEEE Std 1003.1-200x defers to the ISO C standard.

21204       The *isnormal()* macro shall determine whether its argument value is normal (neither zero,  
21205       subnormal, infinite, nor NaN). First, an argument represented in a format wider than its  
21206       semantic type is converted to its semantic type. Then determination is based on the type of the  
21207       argument.

21208 **RETURN VALUE**

21209       The *isnormal()* macro shall return a non-zero value if and only if its argument has a normal  
21210       value.

21211 **ERRORS**

21212       No errors are defined.

21213 **EXAMPLES**

21214       None.

21215 **APPLICATION USAGE**

21216       None.

21217 **RATIONALE**

21218       None.

21219 **FUTURE DIRECTIONS**

21220       None.

21221 **SEE ALSO**

21222       *fpclassify()*, *isfinite()*, *isinf()*, *isnan()*, *signbit()*, the Base Definitions volume of  
21223       IEEE Std 1003.1-200x, <math.h>

21224 **CHANGE HISTORY**

21225       First released in Issue 6. Derived from the ISO/IEC 9899:1999 standard.

21226 **NAME**

21227 isprint — test for a printable character |

21228 **SYNOPSIS**

21229 #include &lt;ctype.h&gt;

21230 int isprint(int c);

21231 **DESCRIPTION**

21232 **CX** The functionality described on this reference page is aligned with the ISO C standard. Any  
21233 conflict between the requirements described here and the ISO C standard is unintentional. This  
21234 volume of IEEE Std 1003.1-200x defers to the ISO C standard.

21235 The *isprint()* function shall test whether *c* is a character of class **print** in the program's current  
21236 locale; see the Base Definitions volume of IEEE Std 1003.1-200x, Chapter 7, Locale.

21237 The *c* argument is an **int**, the value of which the application shall ensure is a character  
21238 representable as an **unsigned char** or equal to the value of the macro EOF. If the argument has  
21239 any other value, the behavior is undefined. |

21240 **RETURN VALUE**

21241 The *isprint()* function shall return non-zero if *c* is a printable character; otherwise, it shall return  
21242 0. |

21243 **ERRORS**

21244 No errors are defined.

21245 **EXAMPLES**

21246 None.

21247 **APPLICATION USAGE**

21248 To ensure applications portability, especially across natural languages, only this function and  
21249 those listed in the SEE ALSO section should be used for character classification.

21250 **RATIONALE**

21251 None.

21252 **FUTURE DIRECTIONS**

21253 None.

21254 **SEE ALSO**

21255 *isalnum()*, *isalpha()*, *iscntrl()*, *isdigit()*, *isgraph()*, *islower()*, *ispunct()*, *isspace()*, *isupper()*,  
21256 *isxdigit()*, *setlocale()*, the Base Definitions volume of IEEE Std 1003.1-200x, <ctype.h>, the Base  
21257 Definitions volume of IEEE Std 1003.1-200x, Chapter 7, Locale

21258 **CHANGE HISTORY**

21259 First released in Issue 1. Derived from Issue 1 of the SVID.

21260 **Issue 6**

21261 The DESCRIPTION is updated to avoid use of the term “must” for application requirements.

21262 **NAME**

21263           ispunct — test for a punctuation character

21264 **SYNOPSIS**

21265           #include &lt;ctype.h&gt;

21266           int ispunct(int c);

21267 **DESCRIPTION**

21268 **CX**       The functionality described on this reference page is aligned with the ISO C standard. Any  
21269 conflict between the requirements described here and the ISO C standard is unintentional. This  
21270 volume of IEEE Std 1003.1-200x defers to the ISO C standard.

21271       The *ispunct()* function shall test whether *c* is a character of class **punct** in the program's current  
21272 locale; see the Base Definitions volume of IEEE Std 1003.1-200x, Chapter 7, Locale.

21273       The *c* argument is an **int**, the value of which the application shall ensure is a character  
21274 representable as an **unsigned char** or equal to the value of the macro EOF. If the argument has  
21275 any other value, the behavior is undefined.

21276 **RETURN VALUE**

21277       The *ispunct()* function shall return non-zero if *c* is a punctuation character; otherwise, it shall  
21278 return 0.

21279 **ERRORS**

21280       No errors are defined.

21281 **EXAMPLES**

21282       None.

21283 **APPLICATION USAGE**

21284       To ensure applications portability, especially across natural languages, only this function and  
21285 those listed in the SEE ALSO section should be used for character classification.

21286 **RATIONALE**

21287       None.

21288 **FUTURE DIRECTIONS**

21289       None.

21290 **SEE ALSO**

21291       *isalnum()*, *isalpha()*, *iscntrl()*, *isdigit()*, *isgraph()*, *islower()*, *isprint()*, *isspace()*, *isupper()*, *isxdigit()*,  
21292 *setlocale()*, the Base Definitions volume of IEEE Std 1003.1-200x, <ctype.h>, the Base Definitions  
21293 volume of IEEE Std 1003.1-200x, Chapter 7, Locale

21294 **CHANGE HISTORY**

21295       First released in Issue 1. Derived from Issue 1 of the SVID.

21296 **Issue 6**

21297       The DESCRIPTION is updated to avoid use of the term “must” for application requirements.

21298 **NAME**

21299           isspace — test for a white-space character

21300 **SYNOPSIS**

21301           #include <ctype.h>

21302           int isspace(int c);

21303 **DESCRIPTION**

21304 **CX**       The functionality described on this reference page is aligned with the ISO C standard. Any  
21305 conflict between the requirements described here and the ISO C standard is unintentional. This  
21306 volume of IEEE Std 1003.1-200x defers to the ISO C standard.

21307       The *isspace()* function shall test whether *c* is a character of class **space** in the program's current  
21308 locale; see the Base Definitions volume of IEEE Std 1003.1-200x, Chapter 7, Locale.

21309       The *c* argument is an **int**, the value of which the application shall ensure is a character  
21310 representable as an **unsigned char** or equal to the value of the macro EOF. If the argument has  
21311 any other value, the behavior is undefined.

21312 **RETURN VALUE**

21313       The *isspace()* function shall return non-zero if *c* is a white-space character; otherwise, it shall  
21314 return 0.

21315 **ERRORS**

21316       No errors are defined.

21317 **EXAMPLES**

21318       None.

21319 **APPLICATION USAGE**

21320       To ensure applications portability, especially across natural languages, only this function and  
21321 those listed in the SEE ALSO section should be used for character classification.

21322 **RATIONALE**

21323       None.

21324 **FUTURE DIRECTIONS**

21325       None.

21326 **SEE ALSO**

21327       *isalnum()*, *isalpha()*, *iscntrl()*, *isdigit()*, *isgraph()*, *islower()*, *isprint()*, *ispunct()*, *isupper()*,  
21328 *isxdigit()*, *setlocale()*, the Base Definitions volume of IEEE Std 1003.1-200x, <ctype.h>, the Base  
21329 Definitions volume of IEEE Std 1003.1-200x, Chapter 7, Locale

21330 **CHANGE HISTORY**

21331       First released in Issue 1. Derived from Issue 1 of the SVID.

21332 **Issue 6**

21333       The DESCRIPTION is updated to avoid use of the term “must” for application requirements.



21334 **NAME**

21335           isunordered — test if arguments are unordered

21336 **SYNOPSIS**

21337           #include <math.h>

21338           int isunordered(real-floating x, real-floating y);

21339 **DESCRIPTION**

21340 cx       The functionality described on this reference page is aligned with the ISO C standard. Any  
21341       conflict between the requirements described here and the ISO C standard is unintentional. This  
21342       volume of IEEE Std 1003.1-200x defers to the ISO C standard.

21343       The *isunordered()* macro shall determine whether its arguments are unordered.

21344 **RETURN VALUE**

21345       Upon successful completion, the *isunordered()* macro shall return 1 if its arguments are  
21346       unordered, and 0 otherwise.

21347       If *x* or *y* is NaN, 0 shall be returned.

21348 **ERRORS**

21349       No errors are defined.

21350 **EXAMPLES**

21351       None.

21352 **APPLICATION USAGE**

21353       The relational and equality operators support the usual mathematical relationships between  
21354       numeric values. For any ordered pair of numeric values, exactly one of the relationships (less,  
21355       greater, and equal) is true. Relational operators may raise the invalid floating-point exception  
21356       when argument values are NaNs. For a NaN and a numeric value, or for two NaNs, just the  
21357       unordered relationship is true. This macro is a quiet (non-floating-point exception raising)  
21358       version of a relational operator. It facilitates writing efficient code that accounts for NaNs  
21359       without suffering the invalid floating-point exception. In the SYNOPSIS section, **real-floating**  
21360       indicates that the argument shall be an expression of **real-floating** type.

21361 **RATIONALE**

21362       None.

21363 **FUTURE DIRECTIONS**

21364       None.

21365 **SEE ALSO**

21366       *isgreater()*, *isgreaterequal()*, *isless()*, *islessequal()*, *islessgreater()*, the Base Definitions volume of  
21367       IEEE Std 1003.1-200x, <math.h>

21368 **CHANGE HISTORY**

21369       First released in Issue 6. Derived from the ISO/IEC 9899:1999 standard.

21370 **NAME**

21371           isupper — test for an uppercase letter

21372 **SYNOPSIS**

21373           #include <ctype.h>

21374           int isupper(int c);

21375 **DESCRIPTION**

21376 **CX**       The functionality described on this reference page is aligned with the ISO C standard. Any  
21377       conflict between the requirements described here and the ISO C standard is unintentional. This  
21378       volume of IEEE Std 1003.1-200x defers to the ISO C standard.

21379       The *isupper()* function shall test whether *c* is a character of class **upper** in the program's current  
21380       locale; see the Base Definitions volume of IEEE Std 1003.1-200x, Chapter 7, Locale.

21381       The *c* argument is an **int**, the value of which the application shall ensure is a character  
21382       representable as an **unsigned char** or equal to the value of the macro EOF. If the argument has  
21383       any other value, the behavior is undefined.

21384 **RETURN VALUE**

21385       The *isupper()* function shall return non-zero if *c* is an uppercase letter; otherwise, it shall return 0.

21386 **ERRORS**

21387       No errors are defined.

21388 **EXAMPLES**

21389       None.

21390 **APPLICATION USAGE**

21391       To ensure applications portability, especially across natural languages, only this function and  
21392       those listed in the SEE ALSO section should be used for character classification.

21393 **RATIONALE**

21394       None.

21395 **FUTURE DIRECTIONS**

21396       None.

21397 **SEE ALSO**

21398       *isalnum()*, *isalpha()*, *iscntrl()*, *isdigit()*, *isgraph()*, *islower()*, *isprint()*, *ispunct()*, *isspace()*, *isxdigit()*,  
21399       *setlocale()*, the Base Definitions volume of IEEE Std 1003.1-200x, <ctype.h>, the Base Definitions  
21400       volume of IEEE Std 1003.1-200x, Chapter 7, Locale

21401 **CHANGE HISTORY**

21402       First released in Issue 1. Derived from Issue 1 of the SVID.

21403 **Issue 6**

21404       The DESCRIPTION is updated to avoid use of the term “must” for application requirements.

21405 **NAME**

21406           iswalnum — test for an alphanumeric wide-character code

21407 **SYNOPSIS**

21408           #include <wctype.h>

21409           int iswalnum(wint\_t wc);

21410 **DESCRIPTION**

21411 cx       The functionality described on this reference page is aligned with the ISO C standard. Any  
21412 conflict between the requirements described here and the ISO C standard is unintentional. This  
21413 volume of IEEE Std 1003.1-200x defers to the ISO C standard.

21414       The *iswalnum()* function shall test whether *wc* is a wide-character code representing a character  
21415 of class **alpha** or **digit** in the program's current locale; see the Base Definitions volume of  
21416 IEEE Std 1003.1-200x, Chapter 7, Locale.

21417       The *wc* argument is a **wint\_t**, the value of which the application shall ensure is a wide-character  
21418 code corresponding to a valid character in the current locale, or equal to the value of the macro  
21419 WEOF. If the argument has any other value, the behavior is undefined.

21420 **RETURN VALUE**

21421       The *iswalnum()* function shall return non-zero if *wc* is an alphanumeric wide-character code;  
21422 otherwise, it shall return 0.

21423 **ERRORS**

21424       No errors are defined.

21425 **EXAMPLES**

21426       None.

21427 **APPLICATION USAGE**

21428       To ensure applications portability, especially across natural languages, only this function and  
21429 those listed in the SEE ALSO section should be used for classification of wide-character codes.

21430 **RATIONALE**

21431       None.

21432 **FUTURE DIRECTIONS**

21433       None.

21434 **SEE ALSO**

21435       *iswalphabet()*, *iswcntrl()*, *iswctype()*, *iswdigit()*, *iswgraph()*, *iswlower()*, *iswprint()*, *iswpunct()*,  
21436 *iswspace()*, *iswupper()*, *iswxdigit()*, *setlocale()*, the Base Definitions volume of  
21437 IEEE Std 1003.1-200x, <wctype.h>, <wchar.h>, <stdio.h>, the Base Definitions volume of  
21438 IEEE Std 1003.1-200x, Chapter 7, Locale

21439 **CHANGE HISTORY**

21440       First released as a World-wide Portability Interface in Issue 4.

21441 **Issue 5**

21442       The following change has been made in this issue for alignment with  
21443 ISO/IEC 9899:1990/Amendment 1:1995 (E):

- 21444       • The SYNOPSIS has been changed to indicate that this function and associated data types are  
21445       now made visible by inclusion of the <wctype.h> header rather than <wchar.h>.

21446 **Issue 6**

21447

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

21448 **NAME**

21449 iswalpha — test for an alphabetic wide-character code

21450 **SYNOPSIS**

21451 #include &lt;wctype.h&gt;

21452 int iswalpha(wint\_t wc);

21453 **DESCRIPTION**

21454 cx The functionality described on this reference page is aligned with the ISO C standard. Any  
 21455 conflict between the requirements described here and the ISO C standard is unintentional. This  
 21456 volume of IEEE Std 1003.1-200x defers to the ISO C standard.

21457 The *iswalpha()* function shall test whether *wc* is a wide-character code representing a character of  
 21458 class **alpha** in the program's current locale; see the Base Definitions volume of  
 21459 IEEE Std 1003.1-200x, Chapter 7, Locale.

21460 The *wc* argument is a **wint\_t**, the value of which the application shall ensure is a wide-character  
 21461 code corresponding to a valid character in the current locale, or equal to the value of the macro  
 21462 WEOF. If the argument has any other value, the behavior is undefined.

21463 **RETURN VALUE**

21464 The *iswalpha()* function shall return non-zero if *wc* is an alphabetic wide-character code;  
 21465 otherwise, it shall return 0.

21466 **ERRORS**

21467 No errors are defined.

21468 **EXAMPLES**

21469 None.

21470 **APPLICATION USAGE**

21471 To ensure applications portability, especially across natural languages, only this function and  
 21472 those listed in the SEE ALSO section should be used for classification of wide-character codes.

21473 **RATIONALE**

21474 None.

21475 **FUTURE DIRECTIONS**

21476 None.

21477 **SEE ALSO**

21478 *iswalnum()*, *iswcntrl()*, *iswctype()*, *iswdigit()*, *iswgraph()*, *iswlower()*, *iswprint()*, *iswpunct()*,  
 21479 *iswspace()*, *iswupper()*, *iswxdigit()*, *setlocale()*, the Base Definitions volume of  
 21480 IEEE Std 1003.1-200x, <wctype.h>, <wchar.h>, <stdio.h>, the Base Definitions volume of  
 21481 IEEE Std 1003.1-200x, Chapter 7, Locale

21482 **CHANGE HISTORY**

21483 First released in Issue 4.

21484 **Issue 5**

21485 The following change has been made in this issue for alignment with  
 21486 ISO/IEC 9899:1990/Amendment 1:1995 (E):

- 21487 • The SYNOPSIS has been changed to indicate that this function and associated data types are  
 21488 now made visible by inclusion of the <wctype.h> header rather than <wchar.h>.

21489 **Issue 6**

21490

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

21491 **NAME**

21492           iswblank — test for a blank wide-character code

21493 **SYNOPSIS**

21494           #include &lt;wctype.h&gt;

21495           int iswblank(wint\_t wc);

21496 **DESCRIPTION**

21497 cx       The functionality described on this reference page is aligned with the ISO C standard. Any  
21498       conflict between the requirements described here and the ISO C standard is unintentional. This  
21499       volume of IEEE Std 1003.1-200x defers to the ISO C standard.

21500       The *iswblank()* function shall test whether *wc* is a wide-character code representing a character of  
21501       class **blank** in the program's current locale; see the Base Definitions volume of  
21502       IEEE Std 1003.1-200x, Chapter 7, Locale.

21503       The *wc* argument is a **wint\_t**, the value of which the application shall ensure is a wide-character  
21504       code corresponding to a valid character in the current locale, or equal to the value of the macro  
21505       WEOF. If the argument has any other value, the behavior is undefined.

21506 **RETURN VALUE**

21507       The *iswblank()* function shall return non-zero if *wc* is a blank wide-character code; otherwise, it  
21508       shall return 0.

21509 **ERRORS**

21510       No errors are defined.

21511 **EXAMPLES**

21512       None.

21513 **APPLICATION USAGE**

21514       To ensure applications portability, especially across natural languages, only this function and  
21515       those listed in the SEE ALSO section should be used for classification of wide-character codes.

21516 **RATIONALE**

21517       None.

21518 **FUTURE DIRECTIONS**

21519       None.

21520 **SEE ALSO**

21521       *iswalnum()*, *iswalpha()*, *iswcntrl()*, *iswctype()*, *iswdigit()*, *iswgraph()*, *iswlower()*, *iswprint()*,  
21522       *iswpunct()*, *iswspace()*, *iswupper()*, *iswxdigit()*, *setlocale()*, the Base Definitions volume of  
21523       IEEE Std 1003.1-200x, <wchar.h>, <wctype.h>, <stdio.h>

21524 **CHANGE HISTORY**

21525       First released in Issue 6. Derived from the ISO/IEC 9899:1999 standard.

21526 **NAME**

21527           iswcntrl — test for a control wide-character code

21528 **SYNOPSIS**

21529           #include <wctype.h>

21530           int iswcntrl(wint\_t wc);

21531 **DESCRIPTION**

21532 **CX**       The functionality described on this reference page is aligned with the ISO C standard. Any  
21533           conflict between the requirements described here and the ISO C standard is unintentional. This  
21534           volume of IEEE Std 1003.1-200x defers to the ISO C standard.

21535           The *iswcntrl()* function shall test whether *wc* is a wide-character code representing a character of  
21536           class **cntrl** in the program's current locale; see the Base Definitions volume of  
21537           IEEE Std 1003.1-200x, Chapter 7, Locale.

21538           The *wc* argument is a **wint\_t**, the value of which the application shall ensure is a wide-character  
21539           code corresponding to a valid character in the current locale, or equal to the value of the macro  
21540           WEOF. If the argument has any other value, the behavior is undefined.

21541 **RETURN VALUE**

21542           The *iswcntrl()* function shall return non-zero if *wc* is a control wide-character code; otherwise, it  
21543           shall return 0.

21544 **ERRORS**

21545           No errors are defined.

21546 **EXAMPLES**

21547           None.

21548 **APPLICATION USAGE**

21549           To ensure applications portability, especially across natural languages, only this function and  
21550           those listed in the SEE ALSO section should be used for classification of wide-character codes.

21551 **RATIONALE**

21552           None.

21553 **FUTURE DIRECTIONS**

21554           None.

21555 **SEE ALSO**

21556           *iswalnum()*, *iswalpha()*, *iswctype()*, *iswdigit()*, *iswgraph()*, *iswlower()*, *iswprint()*, *iswpunct()*,  
21557           *iswspace()*, *iswupper()*, *iswxdigit()*, *setlocale()*, the Base Definitions volume of  
21558           IEEE Std 1003.1-200x, <wctype.h>, <wchar.h>, the Base Definitions volume of  
21559           IEEE Std 1003.1-200x, Chapter 7, Locale

21560 **CHANGE HISTORY**

21561           First released in Issue 4.

21562 **Issue 5**

21563           The following change has been made in this issue for alignment with  
21564           ISO/IEC 9899:1990/Amendment 1:1995 (E):

- 21565           • The SYNOPSIS has been changed to indicate that this function and associated data types are  
21566           now made visible by inclusion of the <wctype.h> header rather than <wchar.h>.



21567 **Issue 6**

21568

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

21569 **NAME**

21570 iswctype — test character for a specified class

21571 **SYNOPSIS**

21572 #include &lt;wctype.h&gt;

21573 int iswctype(wint\_t *wc*, wctype\_t *charclass*);21574 **DESCRIPTION**

21575 **CX** The functionality described on this reference page is aligned with the ISO C standard. Any  
 21576 conflict between the requirements described here and the ISO C standard is unintentional. This  
 21577 volume of IEEE Std 1003.1-200x defers to the ISO C standard.

21578 The *iswctype()* function shall determine whether the wide-character code *wc* has the character  
 21579 class *charclass*, returning true or false. The *iswctype()* function is defined on WEOF and wide-  
 21580 character codes corresponding to the valid character encodings in the current locale. If the *wc*  
 21581 argument is not in the domain of the function, the result is undefined. If the value of *charclass* is  
 21582 invalid (that is, not obtained by a call to *wctype()* or *charclass* is invalidated by a subsequent call  
 21583 to *setlocale()* that has affected category *LC\_CTYPE*) the result is unspecified.

21584 **RETURN VALUE**

21585 The *iswctype()* function shall return non-zero (true) if and only if *wc* has the property described  
 21586 **CX** by *charclass*. If *charclass* is 0, *iswctype()* shall return 0.

21587 **ERRORS**

21588 No errors are defined.

21589 **EXAMPLES**21590 **Testing for a Valid Character**

```
21591 #include <wctype.h>
21592 ...
21593 int yes_or_no;
21594 wint_t wc;
21595 wctype_t valid_class;
21596 ...
21597 if ((valid_class=wctype("vowel")) == (wctype_t)0)
21598     /* Invalid character class. */
21599     yes_or_no=iswctype(wc,valid_class);
```

21600 **APPLICATION USAGE**

21601 The twelve strings "alnum", "alpha", "blank", "cntrl", "digit", "graph", "lower",  
 21602 "print", "punct", "space", "upper", and "xdigit" are reserved for the standard  
 21603 character classes. In the table below, the functions in the left column are equivalent to the  
 21604 functions in the right column.

21605	iswalnum( <i>wc</i> )	iswctype( <i>wc</i> , wctype("alnum"))
21606	iswalpha( <i>wc</i> )	iswctype( <i>wc</i> , wctype("alpha"))
21607	iswblank( <i>wc</i> )	iswctype( <i>wc</i> , wctype("blank"))
21608	iswcntrl( <i>wc</i> )	iswctype( <i>wc</i> , wctype("cntrl"))
21609	iswdigit( <i>wc</i> )	iswctype( <i>wc</i> , wctype("digit"))
21610	iswgraph( <i>wc</i> )	iswctype( <i>wc</i> , wctype("graph"))
21611	iswlower( <i>wc</i> )	iswctype( <i>wc</i> , wctype("lower"))
21612	iswprint( <i>wc</i> )	iswctype( <i>wc</i> , wctype("print"))
21613	iswpunct( <i>wc</i> )	iswctype( <i>wc</i> , wctype("punct"))
21614	iswspace( <i>wc</i> )	iswctype( <i>wc</i> , wctype("space"))

21615            `iswupper(wc)`      `iswctype(wc, wctype("upper"))`  
 21616            `iswxdigit(wc)`    `iswctype(wc, wctype("xdigit"))`

21617 **RATIONALE**

21618            None.

21619 **FUTURE DIRECTIONS**

21620            None.

21621 **SEE ALSO**

21622            `iswalnum()`, `iswalpha()`, `iswcntrl()`, `iswdigit()`, `iswgraph()`, `iswlower()`, `iswprint()`, `iswpunct()`,  
 21623            `iswspace()`, `iswupper()`, `iswxdigit()`, `setlocale()`, `wctype()`, the Base Definitions volume of  
 21624            IEEE Std 1003.1-200x, `<wctype.h>`, `<wchar.h>`

21625 **CHANGE HISTORY**

21626            First released as World-wide Portability Interfaces in Issue 4.

21627 **Issue 5**

21628            The following change has been made in this issue for alignment with  
 21629            ISO/IEC 9899:1990/Amendment 1:1995 (E):

- 21630            • The SYNOPSIS has been changed to indicate that this function and associated data types are  
 21631            now made visible by inclusion of the `<wctype.h>` header rather than `<wchar.h>`.

21632 **Issue 6**

21633            The behavior of  $n=0$  is now described. |

21634            An example is added. |

21635            A new function, `iswblank()`, is added to the list in the APPLICATION USAGE. |

21636 **NAME**

21637 iswdigit — test for a decimal digit wide-character code

21638 **SYNOPSIS**

21639 #include <wctype.h>

21640 int iswdigit(wint\_t wc);

21641 **DESCRIPTION**

21642 cx The functionality described on this reference page is aligned with the ISO C standard. Any  
21643 conflict between the requirements described here and the ISO C standard is unintentional. This  
21644 volume of IEEE Std 1003.1-200x defers to the ISO C standard.

21645 The *iswdigit()* function shall test whether *wc* is a wide-character code representing a character of  
21646 class **digit** in the program's current locale; see the Base Definitions volume of  
21647 IEEE Std 1003.1-200x, Chapter 7, Locale.

21648 The *wc* argument is a **wint\_t**, the value of which the application shall ensure is a wide-character  
21649 code corresponding to a valid character in the current locale, or equal to the value of the macro  
21650 WEOF. If the argument has any other value, the behavior is undefined.

21651 **RETURN VALUE**

21652 The *iswdigit()* function shall return non-zero if *wc* is a decimal digit wide-character code;  
21653 otherwise, it shall return 0.

21654 **ERRORS**

21655 No errors are defined.

21656 **EXAMPLES**

21657 None.

21658 **APPLICATION USAGE**

21659 To ensure applications portability, especially across natural languages, only this function and  
21660 those listed in the SEE ALSO section should be used for classification of wide-character codes.

21661 **RATIONALE**

21662 None.

21663 **FUTURE DIRECTIONS**

21664 None.

21665 **SEE ALSO**

21666 *iswalnum()*, *iswalpha()*, *iswcntrl()*, *iswctype()*, *iswgraph()*, *iswlower()*, *iswprint()*, *iswpunct()*,  
21667 *iswspace()*, *iswupper()*, *iswxdigit()*, *setlocale()*, the Base Definitions volume of  
21668 IEEE Std 1003.1-200x, <wctype.h>, <wchar.h>

21669 **CHANGE HISTORY**

21670 First released in Issue 4.

21671 **Issue 5**

21672 The following change has been made in this issue for alignment with  
21673 ISO/IEC 9899:1990/Amendment 1:1995 (E):

- 21674 • The SYNOPSIS has been changed to indicate that this function and associated data types are  
21675 now made visible by inclusion of the <wctype.h> header rather than <wchar.h>.

21676 **Issue 6**

21677 The DESCRIPTION is updated to avoid use of the term “must” for application requirements.

21678 **NAME**

21679 iswgraph — test for a visible wide-character code

21680 **SYNOPSIS**

21681 #include &lt;wctype.h&gt;

21682 int iswgraph(wint\_t wc);

21683 **DESCRIPTION**

21684 cx The functionality described on this reference page is aligned with the ISO C standard. Any  
 21685 conflict between the requirements described here and the ISO C standard is unintentional. This  
 21686 volume of IEEE Std 1003.1-200x defers to the ISO C standard.

21687 The *iswgraph()* function shall test whether *wc* is a wide-character code representing a character  
 21688 of class **graph** in the program's current locale; see the Base Definitions volume of  
 21689 IEEE Std 1003.1-200x, Chapter 7, Locale.

21690 The *wc* argument is a **wint\_t**, the value of which the application shall ensure is a wide-character  
 21691 code corresponding to a valid character in the current locale, or equal to the value of the macro  
 21692 WEOF. If the argument has any other value, the behavior is undefined.

21693 **RETURN VALUE**

21694 The *iswgraph()* function shall return non-zero if *wc* is a wide-character code with a visible  
 21695 representation; otherwise, it shall return 0.

21696 **ERRORS**

21697 No errors are defined.

21698 **EXAMPLES**

21699 None.

21700 **APPLICATION USAGE**

21701 To ensure applications portability, especially across natural languages, only this function and  
 21702 those listed in the SEE ALSO section should be used for classification of wide-character codes.

21703 **RATIONALE**

21704 None.

21705 **FUTURE DIRECTIONS**

21706 None.

21707 **SEE ALSO**

21708 *iswalnum()*, *iswalpha()*, *iswcntrl()*, *iswctype()*, *iswdigit()*, *iswlower()*, *iswprint()*, *iswpunct()*,  
 21709 *iswspace()*, *iswupper()*, *iswxdigit()*, *setlocale()*, the Base Definitions volume of  
 21710 IEEE Std 1003.1-200x, <**wctype.h**>, <**wchar.h**>, the Base Definitions volume of  
 21711 IEEE Std 1003.1-200x, Chapter 7, Locale

21712 **CHANGE HISTORY**

21713 First released in Issue 4.

21714 **Issue 5**

21715 The following change has been made in this issue for alignment with  
 21716 ISO/IEC 9899:1990/Amendment 1:1995 (E):

- 21717 • The SYNOPSIS has been changed to indicate that this function and associated data types are  
 21718 now made visible by inclusion of the <**wctype.h**> header rather than <**wchar.h**>.

21719 **Issue 6**

21720

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

21721 **NAME**

21722 iswlower — test for a lowercase letter wide-character code

21723 **SYNOPSIS**

21724 #include &lt;wctype.h&gt;

21725 int iswlower(wint\_t wc);

21726 **DESCRIPTION**

21727 *cx* The functionality described on this reference page is aligned with the ISO C standard. Any  
 21728 conflict between the requirements described here and the ISO C standard is unintentional. This  
 21729 volume of IEEE Std 1003.1-200x defers to the ISO C standard.

21730 The *iswlower()* function shall test whether *wc* is a wide-character code representing a character  
 21731 of class **lower** in the program's current locale; see the Base Definitions volume of  
 21732 IEEE Std 1003.1-200x, Chapter 7, Locale.

21733 The *wc* argument is a **wint\_t**, the value of which the application shall ensure is a wide-character  
 21734 code corresponding to a valid character in the current locale, or equal to the value of the macro  
 21735 WEOF. If the argument has any other value, the behavior is undefined.

21736 **RETURN VALUE**

21737 The *iswlower()* function shall return non-zero if *wc* is a lowercase letter wide-character code;  
 21738 otherwise, it shall return 0.

21739 **ERRORS**

21740 No errors are defined.

21741 **EXAMPLES**

21742 None.

21743 **APPLICATION USAGE**

21744 To ensure applications portability, especially across natural languages, only this function and  
 21745 those listed in the SEE ALSO section should be used for classification of wide-character codes.

21746 **RATIONALE**

21747 None.

21748 **FUTURE DIRECTIONS**

21749 None.

21750 **SEE ALSO**

21751 *iswalnum()*, *iswalpha()*, *iswcntrl()*, *iswctype()*, *iswdigit()*, *iswgraph()*, *iswprint()*, *iswpunct()*,  
 21752 *iswspace()*, *iswupper()*, *iswxdigit()*, *setlocale()*, the Base Definitions volume of  
 21753 IEEE Std 1003.1-200x, <wctype.h>, <wchar.h>, the Base Definitions volume of  
 21754 IEEE Std 1003.1-200x, Chapter 7, Locale

21755 **CHANGE HISTORY**

21756 First released in Issue 4.

21757 **Issue 5**

21758 The following change has been made in this issue for alignment with  
 21759 ISO/IEC 9899:1990/Amendment 1:1995 (E):

- 21760 • The SYNOPSIS has been changed to indicate that this function and associated data types are  
 21761 now made visible by inclusion of the <wctype.h> header rather than <wchar.h>.

21762 **Issue 6**

21763

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



21764 **NAME**

21765 iswprint — test for a printable wide-character code |

21766 **SYNOPSIS**

21767 #include &lt;wctype.h&gt;

21768 int iswprint(wint\_t wc);

21769 **DESCRIPTION**21770 cx The functionality described on this reference page is aligned with the ISO C standard. Any  
21771 conflict between the requirements described here and the ISO C standard is unintentional. This  
21772 volume of IEEE Std 1003.1-200x defers to the ISO C standard.21773 The *iswprint()* function shall test whether *wc* is a wide-character code representing a character of  
21774 class **print** in the program's current locale; see the Base Definitions volume of  
21775 IEEE Std 1003.1-200x, Chapter 7, Locale.21776 The *wc* argument is a **wint\_t**, the value of which the application shall ensure is a wide-character |  
21777 code corresponding to a valid character in the current locale, or equal to the value of the macro  
21778 WEOF. If the argument has any other value, the behavior is undefined.21779 **RETURN VALUE**21780 The *iswprint()* function shall return non-zero if *wc* is a printable wide-character code; otherwise, |  
21781 it shall return 0. |21782 **ERRORS**

21783 No errors are defined.

21784 **EXAMPLES**

21785 None.

21786 **APPLICATION USAGE**21787 To ensure applications portability, especially across natural languages, only this function and  
21788 those listed in the SEE ALSO section should be used for classification of wide-character codes.21789 **RATIONALE**

21790 None.

21791 **FUTURE DIRECTIONS**

21792 None.

21793 **SEE ALSO**21794 *iswalnum()*, *iswalpha()*, *iswcntrl()*, *iswctype()*, *iswdigit()*, *iswgraph()*, *iswlower()*, *iswpunct()*,  
21795 *iswspace()*, *iswupper()*, *iswxdigit()*, *setlocale()*, the Base Definitions volume of  
21796 IEEE Std 1003.1-200x, <wctype.h>, <wchar.h>, the Base Definitions volume of  
21797 IEEE Std 1003.1-200x, Chapter 7, Locale21798 **CHANGE HISTORY**

21799 First released in Issue 4.

21800 **Issue 5**21801 The following change has been made in this issue for alignment with  
21802 ISO/IEC 9899:1990/Amendment 1:1995 (E):

- 21803
- The SYNOPSIS has been changed to indicate that this function and associated data types are  
21804 now made visible by inclusion of the <wctype.h> header rather than <wchar.h>.

21805 **Issue 6**

21806

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

21807 **NAME**

21808           iswpunct — test for a punctuation wide-character code

21809 **SYNOPSIS**

21810           #include &lt;wctype.h&gt;

21811           int iswpunct(wint\_t wc);

21812 **DESCRIPTION**

21813 cx       The functionality described on this reference page is aligned with the ISO C standard. Any  
 21814       conflict between the requirements described here and the ISO C standard is unintentional. This  
 21815       volume of IEEE Std 1003.1-200x defers to the ISO C standard.

21816       The *iswpunct()* function shall test whether *wc* is a wide-character code representing a character  
 21817       of class **punct** in the program's current locale; see the Base Definitions volume of  
 21818       IEEE Std 1003.1-200x, Chapter 7, Locale.

21819       The *wc* argument is a **wint\_t**, the value of which the application shall ensure is a wide-character  
 21820       code corresponding to a valid character in the current locale, or equal to the value of the macro  
 21821       WEOF. If the argument has any other value, the behavior is undefined.

21822 **RETURN VALUE**

21823       The *iswpunct()* function shall return non-zero if *wc* is a punctuation wide-character code;  
 21824       otherwise, it shall return 0.

21825 **ERRORS**

21826       No errors are defined.

21827 **EXAMPLES**

21828       None.

21829 **APPLICATION USAGE**

21830       To ensure applications portability, especially across natural languages, only this function and  
 21831       those listed in the SEE ALSO section should be used for classification of wide-character codes.

21832 **RATIONALE**

21833       None.

21834 **FUTURE DIRECTIONS**

21835       None.

21836 **SEE ALSO**

21837       *iswalnum()*, *iswalpha()*, *iswcntrl()*, *iswctype()*, *iswdigit()*, *iswgraph()*, *iswlower()*, *iswprint()*,  
 21838       *iswspace()*, *iswupper()*, *iswxdigit()*, *setlocale()*, the Base Definitions volume of  
 21839       IEEE Std 1003.1-200x, <wctype.h>, <wchar.h>, the Base Definitions volume of  
 21840       IEEE Std 1003.1-200x, Chapter 7, Locale

21841 **CHANGE HISTORY**

21842       First released in Issue 4.

21843 **Issue 5**

21844       The following change has been made in this issue for alignment with  
 21845       ISO/IEC 9899:1990/Amendment 1:1995 (E):

- 21846       • The SYNOPSIS has been changed to indicate that this function and associated data types are  
 21847       now made visible by inclusion of the <wctype.h> header rather than <wchar.h>.

21848 **Issue 6**

21849 The DESCRIPTION is updated to avoid use of the term “must” for application requirements.

21850 **NAME**

21851           iswspace — test for a white-space wide-character code

21852 **SYNOPSIS**

21853           #include <wctype.h>

21854           int iswspace(wint\_t wc);

21855 **DESCRIPTION**

21856 **CX**       The functionality described on this reference page is aligned with the ISO C standard. Any  
21857 conflict between the requirements described here and the ISO C standard is unintentional. This  
21858 volume of IEEE Std 1003.1-200x defers to the ISO C standard.

21859       The *iswspace()* function shall test whether *wc* is a wide-character code representing a character of  
21860 class **space** in the program's current locale; see the Base Definitions volume of  
21861 IEEE Std 1003.1-200x, Chapter 7, Locale.

21862       The *wc* argument is a **wint\_t**, the value of which the application shall ensure is a wide-character  
21863 code corresponding to a valid character in the current locale, or equal to the value of the macro  
21864 WEOF. If the argument has any other value, the behavior is undefined.

21865 **RETURN VALUE**

21866       The *iswspace()* function shall return non-zero if *wc* is a white-space wide-character code;  
21867 otherwise, it shall return 0.

21868 **ERRORS**

21869       No errors are defined.

21870 **EXAMPLES**

21871       None.

21872 **APPLICATION USAGE**

21873       To ensure applications portability, especially across natural languages, only this function and  
21874 those listed in the SEE ALSO section should be used for classification of wide-character codes.

21875 **RATIONALE**

21876       None.

21877 **FUTURE DIRECTIONS**

21878       None.

21879 **SEE ALSO**

21880       *iswalnum()*, *iswalpha()*, *iswcntrl()*, *iswctype()*, *iswdigit()*, *iswgraph()*, *iswlower()*, *iswprint()*,  
21881 *iswpunct()*, *iswupper()*, *iswxdigit()*, *setlocale()*, the Base Definitions volume of  
21882 IEEE Std 1003.1-200x, <**wctype.h**>, <**wchar.h**>, the Base Definitions volume of  
21883 IEEE Std 1003.1-200x, Chapter 7, Locale

21884 **CHANGE HISTORY**

21885       First released in Issue 4.

21886 **Issue 5**

21887       The following change has been made in this issue for alignment with  
21888 ISO/IEC 9899:1990/Amendment 1:1995 (E):

- 21889       • The SYNOPSIS has been changed to indicate that this function and associated data types are  
21890       now made visible by inclusion of the <**wctype.h**> header rather than <**wchar.h**>.

21891 **Issue 6**

21892 The DESCRIPTION is updated to avoid use of the term “must” for application requirements.

21893 **NAME**

21894           iswupper — test for an uppercase letter wide-character code

21895 **SYNOPSIS**

21896           #include <wctype.h>

21897           int iswupper(wint\_t wc);

21898 **DESCRIPTION**

21899 **CX**       The functionality described on this reference page is aligned with the ISO C standard. Any  
21900 conflict between the requirements described here and the ISO C standard is unintentional. This  
21901 volume of IEEE Std 1003.1-200x defers to the ISO C standard.

21902       The *iswupper()* function shall test whether *wc* is a wide-character code representing a character  
21903 of class **upper** in the program's current locale; see the Base Definitions volume of  
21904 IEEE Std 1003.1-200x, Chapter 7, Locale.

21905       The *wc* argument is a **wint\_t**, the value of which the application shall ensure is a wide-character  
21906 code corresponding to a valid character in the current locale, or equal to the value of the macro  
21907 WEOF. If the argument has any other value, the behavior is undefined.

21908 **RETURN VALUE**

21909       The *iswupper()* function shall return non-zero if *wc* is an uppercase letter wide-character code;  
21910 otherwise, it shall return 0.

21911 **ERRORS**

21912       No errors are defined.

21913 **EXAMPLES**

21914       None.

21915 **APPLICATION USAGE**

21916       To ensure applications portability, especially across natural languages, only this function and  
21917 those listed in the SEE ALSO section should be used for classification of wide-character codes.

21918 **RATIONALE**

21919       None.

21920 **FUTURE DIRECTIONS**

21921       None.

21922 **SEE ALSO**

21923       *iswalnum()*, *iswalpha()*, *iswcntrl()*, *iswctype()*, *iswdigit()*, *iswgraph()*, *iswlower()*, *iswprint()*,  
21924 *iswpunct()*, *iswspace()*, *iswxdigit()*, *setlocale()*, the Base Definitions volume of  
21925 IEEE Std 1003.1-200x, <**wctype.h**>, <**wchar.h**>, the Base Definitions volume of  
21926 IEEE Std 1003.1-200x, Chapter 7, Locale

21927 **CHANGE HISTORY**

21928       First released in Issue 4.

21929 **Issue 5**

21930       The following change has been made in this issue for alignment with  
21931 ISO/IEC 9899:1990/Amendment 1:1995 (E):

- 21932       • The SYNOPSIS has been changed to indicate that this function and associated data types are  
21933       now made visible by inclusion of the <**wctype.h**> header rather than <**wchar.h**>.

21934 **Issue 6**

21935 The DESCRIPTION is updated to avoid use of the term “must” for application requirements.



21936 **NAME**

21937 iswxdigit — test for a hexadecimal digit wide-character code

21938 **SYNOPSIS**

21939 #include &lt;wctype.h&gt;

21940 int iswxdigit(wint\_t wc);

21941 **DESCRIPTION**

21942 cx The functionality described on this reference page is aligned with the ISO C standard. Any  
 21943 conflict between the requirements described here and the ISO C standard is unintentional. This  
 21944 volume of IEEE Std 1003.1-200x defers to the ISO C standard.

21945 The *iswxdigit()* function shall test whether *wc* is a wide-character code representing a character  
 21946 of class **xdigit** in the program's current locale; see the Base Definitions volume of  
 21947 IEEE Std 1003.1-200x, Chapter 7, Locale.

21948 The *wc* argument is a **wint\_t**, the value of which the application shall ensure is a wide-character  
 21949 code corresponding to a valid character in the current locale, or equal to the value of the macro  
 21950 WEOF. If the argument has any other value, the behavior is undefined.

21951 **RETURN VALUE**

21952 The *iswxdigit()* function shall return non-zero if *wc* is a hexadecimal digit wide-character code;  
 21953 otherwise, it shall return 0.

21954 **ERRORS**

21955 No errors are defined.

21956 **EXAMPLES**

21957 None.

21958 **APPLICATION USAGE**

21959 To ensure applications portability, especially across natural languages, only this function and  
 21960 those listed in the SEE ALSO section should be used for classification of wide-character codes.

21961 **RATIONALE**

21962 None.

21963 **FUTURE DIRECTIONS**

21964 None.

21965 **SEE ALSO**

21966 *iswalnum()*, *iswalpha()*, *iswcntrl()*, *iswctype()*, *iswdigit()*, *iswgraph()*, *iswlower()*, *iswprint()*,  
 21967 *iswpunct()*, *iswspace()*, *iswupper()*, *setlocale()*, the Base Definitions volume of  
 21968 IEEE Std 1003.1-200x, <wctype.h>, <wchar.h>

21969 **CHANGE HISTORY**

21970 First released in Issue 4.

21971 **Issue 5**

21972 The following change has been made in this issue for alignment with  
 21973 ISO/IEC 9899:1990/Amendment 1:1995 (E):

- 21974 • The SYNOPSIS has been changed to indicate that this function and associated data types are  
 21975 now made visible by inclusion of the <wctype.h> header rather than <wchar.h>.

21976 **Issue 6**

21977 The DESCRIPTION is updated to avoid use of the term “must” for application requirements.

21978 **NAME**

21979 isxdigit — test for a hexadecimal digit

21980 **SYNOPSIS**

21981 #include <ctype.h>

21982 int isxdigit(int c);

21983 **DESCRIPTION**

21984 cx The functionality described on this reference page is aligned with the ISO C standard. Any  
21985 conflict between the requirements described here and the ISO C standard is unintentional. This  
21986 volume of IEEE Std 1003.1-200x defers to the ISO C standard.

21987 The *isxdigit()* function shall test whether *c* is a character of class **xdigit** in the program's current  
21988 locale; see the Base Definitions volume of IEEE Std 1003.1-200x, Chapter 7, Locale.

21989 The *c* argument is an **int**, the value of which the application shall ensure is a character  
21990 representable as an **unsigned char** or equal to the value of the macro EOF. If the argument has  
21991 any other value, the behavior is undefined.

21992 **RETURN VALUE**

21993 The *isxdigit()* function shall return non-zero if *c* is a hexadecimal digit; otherwise, it shall return  
21994 0.

21995 **ERRORS**

21996 No errors are defined.

21997 **EXAMPLES**

21998 None.

21999 **APPLICATION USAGE**

22000 To ensure applications portability, especially across natural languages, only this function and  
22001 those listed in the SEE ALSO section should be used for character classification.

22002 **RATIONALE**

22003 None.

22004 **FUTURE DIRECTIONS**

22005 None.

22006 **SEE ALSO**

22007 *isalnum()*, *isalpha()*, *iscntrl()*, *isdigit()*, *isgraph()*, *islower()*, *isprint()*, *ispunct()*, *isspace()*, *isupper()*,  
22008 the Base Definitions volume of IEEE Std 1003.1-200x, <ctype.h>

22009 **CHANGE HISTORY**

22010 First released in Issue 1. Derived from Issue 1 of the SVID.

22011 **Issue 6**

22012 The DESCRIPTION is updated to avoid use of the term “must” for application requirements.

22013 **NAME**

22014 j0, j1, jn — Bessel functions of the first kind

22015 **SYNOPSIS**

```
22016 xSI #include <math.h>
22017 double j0(double x);
22018 double j1(double x);
22019 double jn(int n, double x);
22020
```

22021 **DESCRIPTION**

22022 The *j0()*, *j1()*, and *jn()* functions shall compute Bessel functions of *x* of the first kind of orders 0,  
 22023 1, and *n*, respectively.

22024 An application wishing to check for error situations should set *errno* to zero and call  
 22025 *feclearexcept*(FE\_ALL\_EXCEPT) before calling these functions. On return, if *errno* is non-zero or  
 22026 *fetestexcept*(FE\_INVALID | FE\_DIVBYZERO | FE\_OVERFLOW | FE\_UNDERFLOW) is non-  
 22027 zero, an error has occurred.

22028 **RETURN VALUE**

22029 Upon successful completion, these functions shall return the relevant Bessel value of *x* of the  
 22030 first kind.

22031 If the *x* argument is too large in magnitude, or the correct result would cause underflow, 0 shall  
 22032 be returned and a range error may occur.

22033 If *x* is NaN, a NaN shall be returned.

22034 **ERRORS**

22035 These functions may fail if:

22036 Range Error The value of *x* was too large in magnitude, or an underflow occurred.

22037 If the integer expression (math\_errhandling & MATH\_ERRNO) is non-zero, |  
 22038 then *errno* shall be set to [ERANGE]. If the integer expression |  
 22039 (math\_errhandling & MATH\_ERREXCEPT) is non-zero, then the underflow |  
 22040 floating-point exception shall be raised. |

22041 No other errors shall occur.

22042 **EXAMPLES**

22043 None.

22044 **APPLICATION USAGE**

22045 On error, the expressions (math\_errhandling & MATH\_ERRNO) and (math\_errhandling &  
 22046 MATH\_ERREXCEPT) are independent of each other, but at least one of them must be non-zero.

22047 **RATIONALE**

22048 None.

22049 **FUTURE DIRECTIONS**

22050 None.

22051 **SEE ALSO**

22052 *feclearexcept()*, *fetestexcept()*, *isnan()*, *y0()*, the Base Definitions volume of IEEE Std 1003.1-200x, |  
 22053 Section 4.18, Treatment of Error Conditions for Mathematical Functions, <math.h> |

22054 **CHANGE HISTORY**

22055 First released in Issue 1. Derived from Issue 1 of the SVID.

22056 **Issue 5**

22057 The DESCRIPTION is updated to indicate how an application should check for an error. This  
22058 text was previously published in the APPLICATION USAGE section.

22059 **Issue 6**

22060 The may fail [EDOM] error is removed for the case for NaN.

22061 The RETURN VALUE and ERRORS sections are reworked for alignment of the error handling |  
22062 with the ISO/IEC 9899:1999 standard.

22063 **NAME**

22064       jrand48 — generate a uniformly distributed pseudo-random long signed integer

22065 **SYNOPSIS**

22066 xSI     #include <stdlib.h>

22067       long jrand48(unsigned short xsubi[3]);

22068

22069 **DESCRIPTION**

22070       Refer to *drand48()*.

## 22071 NAME

22072 kill — send a signal to a process or a group of processes

## 22073 SYNOPSIS

22074 cx #include &lt;signal.h&gt;

22075 int kill(pid\_t pid, int sig);

22076

## 22077 DESCRIPTION

22078 The *kill()* function shall send a signal to a process or a group of processes specified by *pid*. The  
 22079 signal to be sent is specified by *sig* and is either one from the list given in <signal.h> or 0. If *sig* is  
 22080 0 (the null signal), error checking is performed but no signal is actually sent. The null signal can  
 22081 be used to check the validity of *pid*.

22082 For a process to have permission to send a signal to a process designated by *pid*, unless the  
 22083 sending process has appropriate privileges, the real or effective user ID of the sending process  
 22084 shall match the real or saved set-user-ID of the receiving process.

22085 If *pid* is greater than 0, *sig* shall be sent to the process whose process ID is equal to *pid*.

22086 If *pid* is 0, *sig* shall be sent to all processes (excluding an unspecified set of system processes)  
 22087 whose process group ID is equal to the process group ID of the sender, and for which the  
 22088 process has permission to send a signal.

22089 If *pid* is  $-1$ , *sig* shall be sent to all processes (excluding an unspecified set of system processes) for  
 22090 which the process has permission to send that signal.

22091 If *pid* is negative, but not  $-1$ , *sig* shall be sent to all processes (excluding an unspecified set of  
 22092 system processes) whose process group ID is equal to the absolute value of *pid*, and for which  
 22093 the process has permission to send a signal.

22094 If the value of *pid* causes *sig* to be generated for the sending process, and if *sig* is not blocked for  
 22095 the calling thread and if no other thread has *sig* unblocked or is waiting in a *sigwait()* function  
 22096 for *sig*, either *sig* or at least one pending unblocked signal shall be delivered to the sending  
 22097 thread before *kill()* returns.

22098 The user ID tests described above shall not be applied when sending SIGCONT to a process that  
 22099 is a member of the same session as the sending process.

22100 An implementation that provides extended security controls may impose further  
 22101 implementation-defined restrictions on the sending of signals, including the null signal. In  
 22102 particular, the system may deny the existence of some or all of the processes specified by *pid*.

22103 The *kill()* function is successful if the process has permission to send *sig* to any of the processes  
 22104 specified by *pid*. If *kill()* fails, no signal shall be sent.

## 22105 RETURN VALUE

22106 Upon successful completion, 0 shall be returned. Otherwise,  $-1$  shall be returned and *errno* set to  
 22107 indicate the error.

## 22108 ERRORS

22109 The *kill()* function shall fail if:

22110 [EINVAL] The value of the *sig* argument is an invalid or unsupported signal number.

22111 [EPERM] The process does not have permission to send the signal to any receiving  
 22112 process.

22113 [ESRCH] No process or process group can be found corresponding to that specified by  
 22114 *pid*.

22115 **EXAMPLES**

22116 None.

22117 **APPLICATION USAGE**

22118 None.

22119 **RATIONALE**

22120 The semantics for permission checking for *kill()* differed between System V and most other  
 22121 implementations, such as Version 7 or 4.3 BSD. The semantics chosen for this volume of  
 22122 IEEE Std 1003.1-200x agree with System V. Specifically, a set-user-ID process cannot protect  
 22123 itself against signals (or at least not against SIGKILL) unless it changes its real user ID. This  
 22124 choice allows the user who starts an application to send it signals even if it changes its effective  
 22125 user ID. The other semantics give more power to an application that wants to protect itself from  
 22126 the user who ran it.

22127 Some implementations provide semantic extensions to the *kill()* function when the absolute  
 22128 value of *pid* is greater than some maximum, or otherwise special, value. Negative values are a  
 22129 flag to *kill()*. Since most implementations return [ESRCH] in this case, this behavior is not  
 22130 included in this volume of IEEE Std 1003.1-200x, although a conforming implementation could  
 22131 provide such an extension.

22132 The implementation-defined processes to which a signal cannot be sent may include the  
 22133 scheduler or *init*.

22134 There was initially strong sentiment to specify that, if *pid* specifies that a signal be sent to the  
 22135 calling process and that signal is not blocked, that signal would be delivered before *kill()*  
 22136 returns. This would permit a process to call *kill()* and be guaranteed that the call never return.  
 22137 However, historical implementations that provide only the *signal()* function make only the  
 22138 weaker guarantee in this volume of IEEE Std 1003.1-200x, because they only deliver one signal  
 22139 each time a process enters the kernel. Modifications to such implementations to support the  
 22140 *sigaction()* function generally require entry to the kernel following return from a signal-catching  
 22141 function, in order to restore the signal mask. Such modifications have the effect of satisfying the  
 22142 stronger requirement, at least when *sigaction()* is used, but not necessarily when *signal()* is used.  
 22143 The developers of this volume of IEEE Std 1003.1-200x considered making the stronger  
 22144 requirement except when *signal()* is used, but felt this would be unnecessarily complex.  
 22145 Implementors are encouraged to meet the stronger requirement whenever possible. In practice,  
 22146 the weaker requirement is the same, except in the rare case when two signals arrive during a  
 22147 very short window. This reasoning also applies to a similar requirement for *sigprocmask()*.

22148 In 4.2 BSD, the SIGCONT signal can be sent to any descendant process regardless of user-ID  
 22149 security checks. This allows a job control shell to continue a job even if processes in the job have  
 22150 altered their user IDs (as in the *su* command). In keeping with the addition of the concept of  
 22151 sessions, similar functionality is provided by allowing the SIGCONT signal to be sent to any  
 22152 process in the same session regardless of user ID security checks. This is less restrictive than BSD  
 22153 in the sense that ancestor processes (in the same session) can now be the recipient. It is more  
 22154 restrictive than BSD in the sense that descendant processes that form new sessions are now  
 22155 subject to the user ID checks. A similar relaxation of security is not necessary for the other job  
 22156 control signals since those signals are typically sent by the terminal driver in recognition of  
 22157 special characters being typed; the terminal driver bypasses all security checks.

22158 In secure implementations, a process may be restricted from sending a signal to a process having  
 22159 a different security label. In order to prevent the existence or nonexistence of a process from  
 22160 being used as a covert channel, such processes should appear nonexistent to the sender; that is,  
 22161 [ESRCH] should be returned, rather than [EPERM], if *pid* refers only to such processes.

22162 Existing implementations vary on the result of a *kill()* with *pid* indicating an inactive process (a  
 22163 terminated process that has not been waited for by its parent). Some indicate success on such a  
 22164 call (subject to permission checking), while others give an error of [ESRCH]. Since the definition  
 22165 of process lifetime in this volume of IEEE Std 1003.1-200x covers inactive processes, the  
 22166 [ESRCH] error as described is inappropriate in this case. In particular, this means that an  
 22167 application cannot have a parent process check for termination of a particular child with *kill()*.  
 22168 (Usually this is done with the null signal; this can be done reliably with *waitpid()*.)

22169 There is some belief that the name *kill()* is misleading, since the function is not always intended  
 22170 to cause process termination. However, the name is common to all historical implementations,  
 22171 and any change would be in conflict with the goal of minimal changes to existing application  
 22172 code.

#### 22173 FUTURE DIRECTIONS

22174 None.

#### 22175 SEE ALSO

22176 *getpid()*, *raise()*, *setsid()*, *sigaction()*, *sigqueue()*, the Base Definitions volume of  
 22177 IEEE Std 1003.1-200x, <signal.h>, <sys/types.h>

#### 22178 CHANGE HISTORY

22179 First released in Issue 1. Derived from Issue 1 of the SVID.

#### 22180 Issue 5

22181 The DESCRIPTION is updated for alignment with POSIX Threads Extension.

#### 22182 Issue 6

22183 In the SYNOPSIS, the optional include of the <sys/types.h> header is removed.

22184 The following new requirements on POSIX implementations derive from alignment with the  
 22185 Single UNIX Specification:

- 22186 • In the DESCRIPTION, the second paragraph is reworded to indicate that the saved set-user-  
 22187 ID of the calling process is checked in place of its effective user ID. This is a FIPS  
 22188 requirement.
- 22189 • The requirement to include <sys/types.h> has been removed. Although <sys/types.h> was  
 22190 required for conforming implementations of previous POSIX specifications, it was not  
 22191 required for UNIX applications.
- 22192 • The behavior when *pid* is  $-1$  is now specified. It was previously explicitly unspecified in the  
 22193 POSIX.1-1988 standard.

22194 The DESCRIPTION is updated to avoid use of the term “must” for application requirements.



22195 **NAME**

22196 killpg — send a signal to a process group

22197 **SYNOPSIS**

22198 XSI #include &lt;signal.h&gt;

22199 int killpg(pid\_t pgrp, int sig);

22200

22201 **DESCRIPTION**22202 The *killpg()* function shall send the signal specified by *sig* to the process group specified by *pgrp*.22203 If *pgrp* is greater than 1, *killpg(pgrp, sig)* shall be equivalent to *kill(-pgrp, sig)*. If *pgrp* is less than or  
22204 equal to 1, the behavior of *killpg()* is undefined.22205 **RETURN VALUE**22206 Refer to *kill()*.22207 **ERRORS**22208 Refer to *kill()*.22209 **EXAMPLES**

22210 None.

22211 **APPLICATION USAGE**

22212 None.

22213 **RATIONALE**

22214 None.

22215 **FUTURE DIRECTIONS**

22216 None.

22217 **SEE ALSO**22218 *getpgid()*, *getpid()*, *kill()*, *raise()*, the Base Definitions volume of IEEE Std 1003.1-200x, <signal.h>22219 **CHANGE HISTORY**

22220 First released in Issue 4, Version 2.

22221 **Issue 5**

22222 Moved from X/OPEN UNIX extension to BASE.

22223 **NAME**

22224        **l64a** — convert a 32-bit integer to a radix-64 ASCII string

22225 **SYNOPSIS**

22226 XSI     #include <stdlib.h>

22227        char \*l64a(long value);

22228

22229 **DESCRIPTION**

22230        Refer to *a64l()*.

22231 **NAME**

22232 labs, llabs — return a long integer absolute value

22233 **SYNOPSIS**

22234 #include <stdlib.h>

22235 long labs(long *i*);

22236 long long llabs(long long *i*);

22237 **DESCRIPTION**

22238 **CX** The functionality described on this reference page is aligned with the ISO C standard. Any  
22239 conflict between the requirements described here and the ISO C standard is unintentional. This  
22240 volume of IEEE Std 1003.1-200x defers to the ISO C standard.

22241 The *labs()* function shall compute the absolute value of the **long** integer operand *i*. The *llabs()*  
22242 function shall compute the absolute value of the **long long** integer operand *i*. If the result cannot  
22243 be represented, the behavior is undefined.

22244 **RETURN VALUE**

22245 The *labs()* function shall return the absolute value of the **long** integer operand. The *llabs()*  
22246 function shall return the absolute value of the **long long** integer operand.

22247 **ERRORS**

22248 No errors are defined.

22249 **EXAMPLES**

22250 None.

22251 **APPLICATION USAGE**

22252 None.

22253 **RATIONALE**

22254 None.

22255 **FUTURE DIRECTIONS**

22256 None.

22257 **SEE ALSO**

22258 *abs()*, the Base Definitions volume of IEEE Std 1003.1-200x, <stdlib.h>

22259 **CHANGE HISTORY**

22260 First released in Issue 4. Derived from the ISO C standard.

22261 **Issue 6**

22262 The *llabs()* function is added for alignment with the ISO/IEC 9899:1999 standard.

22263 **NAME**

22264 lchown — change the owner and group of a symbolic link

22265 **SYNOPSIS**

22266 XSI #include &lt;unistd.h&gt;

22267 int lchown(const char \*path, uid\_t owner, gid\_t group);

22268

22269 **DESCRIPTION**

22270 The *lchown()* function shall be equivalent to *chown()*, except in the case where the named file is a  
 22271 symbolic link. In this case, *lchown()* shall change the ownership of the symbolic link file itself,  
 22272 while *chown()* changes the ownership of the file or directory to which the symbolic link refers.

22273 **RETURN VALUE**

22274 Upon successful completion, *lchown()* shall return 0. Otherwise, it shall return -1 and set *errno* to  
 22275 indicate an error.

22276 **ERRORS**22277 The *lchown()* function shall fail if:22278 [EACCES] Search permission is denied on a component of the path prefix of *path*.

22279 [EINVAL] The owner or group ID is not a value supported by the implementation.

22280 [ELOOP] A loop exists in symbolic links encountered during resolution of the *path*  
 22281 argument.

22282 [ENAMETOOLONG]

22283 The length of a pathname exceeds {PATH\_MAX} or a pathname component is  
 22284 longer than {NAME\_MAX}.

22285 [ENOENT] A component of *path* does not name an existing file or *path* is an empty string.22286 [ENOTDIR] A component of the path prefix of *path* is not a directory.

22287 [EOPNOTSUPP] The *path* argument names a symbolic link and the implementation does not  
 22288 support setting the owner or group of a symbolic link.

22289 [EPERM] The effective user ID does not match the owner of the file and the process  
 22290 does not have appropriate privileges.

22291 [EROFS] The file resides on a read-only file system.

22292 The *lchown()* function may fail if:

22293 [EIO] An I/O error occurred while reading or writing to the file system.

22294 [EINTR] A signal was caught during execution of the function.

22295 [ELOOP] More than {SYMLOOP\_MAX} symbolic links were encountered during  
 22296 resolution of the *path* argument.

22297 [ENAMETOOLONG]

22298 Pathname resolution of a symbolic link produced an intermediate result  
 22299 whose length exceeds {PATH\_MAX}.

22300 **EXAMPLES**22301 **Changing the Current Owner of a File**

22302 The following example shows how to change the ownership of the symbolic link named  
22303 **/modules/pass1** to the user ID associated with “jones” and the group ID associated with “cnd”.

22304 The numeric value for the user ID is obtained by using the *getpwnam()* function. The numeric  
22305 value for the group ID is obtained by using the *getgrnam()* function.

```
22306 #include <sys/types.h>
22307 #include <unistd.h>
22308 #include <pwd.h>
22309 #include <grp.h>

22310 struct passwd *pwd;
22311 struct group *grp;
22312 char *path = "/modules/pass1";
22313 ...
22314 pwd = getpwnam("jones");
22315 grp = getgrnam("cnd");
22316 lchown(path, pwd->pw_uid, grp->gr_gid);
```

22317 **APPLICATION USAGE**

22318 None.

22319 **RATIONALE**

22320 None.

22321 **FUTURE DIRECTIONS**

22322 None.

22323 **SEE ALSO**

22324 *chown()*, *symlink()*, the Base Definitions volume of IEEE Std 1003.1-200x, **<unistd.h>**

22325 **CHANGE HISTORY**

22326 First released in Issue 4, Version 2.

22327 **Issue 5**

22328 Moved from X/OPEN UNIX extension to BASE.

22329 **Issue 6**

22330 The wording of the mandatory [ELOOP] error condition is updated, and a second optional  
22331 [ELOOP] error condition is added.

22332 **NAME**

22333        lcong48 — seed a uniformly distributed pseudo-random signed long integer generator

22334 **SYNOPSIS**

22335 xSI     #include <stdlib.h>

22336        void lcong48(unsigned short param[7]);

22337

22338 **DESCRIPTION**

22339        Refer to *drand48()*.

22340 **NAME**

22341 ldexp, ldexpf, ldexpl — load exponent of a floating-point number

22342 **SYNOPSIS**

22343 #include &lt;math.h&gt;

22344 double ldexp(double x, int exp);

22345 float ldexpf(float x, int exp);

22346 long double ldexpl(long double x, int exp);

22347 **DESCRIPTION**

22348 **CX** The functionality described on this reference page is aligned with the ISO C standard. Any  
 22349 conflict between the requirements described here and the ISO C standard is unintentional. This  
 22350 volume of IEEE Std 1003.1-200x defers to the ISO C standard.

22351 These functions shall compute the quantity  $x * 2^{exp}$ .

22352 An application wishing to check for error situations should set *errno* to zero and call  
 22353 *feclearexcept*(FE\_ALL\_EXCEPT) before calling these functions. On return, if *errno* is non-zero or  
 22354 *fetestexcept*(FE\_INVALID | FE\_DIVBYZERO | FE\_OVERFLOW | FE\_UNDERFLOW) is non-  
 22355 zero, an error has occurred.

22356 **RETURN VALUE**22357 Upon successful completion, these functions shall return *x* multiplied by 2, raised to the power  
22358 *exp*.

22359 If these functions would cause overflow, a range error shall occur and *ldexp*(*x*), *ldexpf*(*x*), and  
 22360 *ldexpl*(*x*) shall return ±HUGE\_VAL, ±HUGE\_VALF, and ±HUGE\_VALL (according to the sign of  
 22361 *x*), respectively.

22362 If the correct value would cause underflow, and is not representable, a range error may occur,  
 22363 **MX** and either 0.0 (if supported), or an implementation-defined value shall be returned.

22364 **MX** If *x* is NaN, a NaN shall be returned.22365 If *x* is ±0 or ±Inf, *x* shall be returned.22366 If *exp* is 0, *x* shall be returned.

22367 If the correct value would cause underflow, and is representable, a range error may occur and  
 22368 the correct value shall be returned.

22369 **ERRORS**

22370 These functions shall fail if:

22371 Range Error The result overflows.

22372 If the integer expression (math\_errhandling & MATH\_ERRNO) is non-zero, |  
 22373 then *errno* shall be set to [ERANGE]. If the integer expression |  
 22374 (math\_errhandling & MATH\_ERREXCEPT) is non-zero, then the overflow |  
 22375 floating-point exception shall be raised. |

22376 These functions may fail if:

22377 Range Error The result underflows.

22378 If the integer expression (math\_errhandling & MATH\_ERRNO) is non-zero, |  
 22379 then *errno* shall be set to [ERANGE]. If the integer expression |  
 22380 (math\_errhandling & MATH\_ERREXCEPT) is non-zero, then the underflow |  
 22381 floating-point exception shall be raised. |

22382 **EXAMPLES**

22383       None.

22384 **APPLICATION USAGE**

22385       On error, the expressions (math\_errhandling & MATH\_ERRNO) and (math\_errhandling &  
22386       MATH\_ERREXCEPT) are independent of each other, but at least one of them must be non-zero.

22387 **RATIONALE**

22388       None.

22389 **FUTURE DIRECTIONS**

22390       None.

22391 **SEE ALSO**

22392       *feclearexcept()*, *fetestexcept()*, *frexp()*, *isnan()*, the Base Definitions volume of |  
22393       IEEE Std 1003.1-200x, Section 4.18, Treatment of Error Conditions for Mathematical Functions, |  
22394       <math.h>

22395 **CHANGE HISTORY**

22396       First released in Issue 1. Derived from Issue 1 of the SVID.

22397 **Issue 5**

22398       The DESCRIPTION is updated to indicate how an application should check for an error. This  
22399       text was previously published in the APPLICATION USAGE section.

22400 **Issue 6**

22401       The *ldexpf()* and *ldexpl()* functions are added for alignment with the ISO/IEC 9899:1999  
22402       standard.

22403       The DESCRIPTION, RETURN VALUE, ERRORS, and APPLICATION USAGE sections are  
22404       revised to align with the ISO/IEC 9899:1999 standard.

22405       IEC 60559:1989 standard floating-point extensions over the ISO/IEC 9899:1999 standard are  
22406       marked.



22407 **NAME**

22408 ldiv, lldiv — compute quotient and remainder of a long division

22409 **SYNOPSIS**

22410 #include &lt;stdlib.h&gt;

22411 ldiv\_t ldiv(long *numer*, long *denom*);22412 lldiv\_t lldiv(long long *numer*, long long *denom*);22413 **DESCRIPTION**

22414 **CX** The functionality described on this reference page is aligned with the ISO C standard. Any  
 22415 conflict between the requirements described here and the ISO C standard is unintentional. This  
 22416 volume of IEEE Std 1003.1-200x defers to the ISO C standard.

22417 These functions shall compute the quotient and remainder of the division of the numerator  
 22418 *numer* by the denominator *denom*. If the division is inexact, the resulting quotient is the **long**  
 22419 integer (for the *ldiv()* function) or **long long** integer (for the *lldiv()* function) of lesser magnitude  
 22420 that is the nearest to the algebraic quotient. If the result cannot be represented, the behavior is  
 22421 undefined; otherwise, *quot* \* *denom* + *rem* shall equal *numer*.

22422 **RETURN VALUE**

22423 The *ldiv()* function shall return a structure of type **ldiv\_t**, comprising both the quotient and the  
 22424 remainder. The structure shall include the following members, in any order:

22425 long quot; /\* Quotient \*/

22426 long rem; /\* Remainder \*/

22427 The *lldiv()* function shall return a structure of type **lldiv\_t**, comprising both the quotient and the  
 22428 remainder. The structure shall include the following members, in any order:

22429 long long quot; /\* Quotient \*/

22430 long long rem; /\* Remainder \*/

22431 **ERRORS**

22432 No errors are defined.

22433 **EXAMPLES**

22434 None.

22435 **APPLICATION USAGE**

22436 None.

22437 **RATIONALE**

22438 None.

22439 **FUTURE DIRECTIONS**

22440 None.

22441 **SEE ALSO**22442 *div()*, the Base Definitions volume of IEEE Std 1003.1-200x, <stdlib.h>22443 **CHANGE HISTORY**

22444 First released in Issue 4. Derived from the ISO C standard.

22445 **Issue 6**22446 The *lldiv()* function is added for alignment with the ISO/IEC 9899:1999 standard.

22447 **NAME**

22448       lfind — find entry in a linear search table

22449 **SYNOPSIS**

22450 XSI       #include <search.h>

```
22451       void *lfind(const void *key, const void *base, size_t *nelp,  
22452                   size_t width, int (*compar)(const void *, const void *));
```

22453

22454 **DESCRIPTION**

22455       Refer to *lsearch()*.

22456 **NAME**

22457 lgamma, lgammaf, lgammal — log gamma function

22458 **SYNOPSIS**

22459 #include &lt;math.h&gt;

22460 double lgamma(double x);

22461 float lgammaf(float x);

22462 long double lgammal(long double x);

22463 XSI extern int signgam;

22464

22465 **DESCRIPTION**

22466 CX The functionality described on this reference page is aligned with the ISO C standard. Any  
 22467 conflict between the requirements described here and the ISO C standard is unintentional. This  
 22468 volume of IEEE Std 1003.1-200x defers to the ISO C standard.

22469 These functions shall compute  $\log_e |\Gamma(x)|$  where  $\Gamma(x)$  is defined as  $\int_0^{\infty} e^{-t} t^{x-1} dt$ . The argument  $x$   
 22470 need not be a non-positive integer ( $\Gamma(x)$  is defined over the reals, except the non-positive  
 22471 integers).  
 22472

22473 XSI The sign of  $\Gamma(x)$  is returned in the external integer *signgam*.

22474 CX These functions need not be reentrant. A function that is not required to be reentrant is not  
 22475 required to be thread-safe.

22476 An application wishing to check for error situations should set *errno* to zero and call  
 22477 *feclearexcept*(FE\_ALL\_EXCEPT) before calling these functions. On return, if *errno* is non-zero or  
 22478 *fetestexcept*(FE\_INVALID | FE\_DIVBYZERO | FE\_OVERFLOW | FE\_UNDERFLOW) is non-  
 22479 zero, an error has occurred.

22480 **RETURN VALUE**22481 Upon successful completion, these functions shall return the logarithmic gamma of  $x$ .

22482 If  $x$  is a non-positive integer, a pole error shall occur and *lgamma()*, *lgammaf()*, and *lgammal()*  
 22483 shall return +HUGE\_VAL, +HUGE\_VALF, and +HUGE\_VALL, respectively.

22484 If the correct value would cause overflow, a range error shall occur and *lgamma()*, *lgammaf()*,  
 22485 and *lgammal()* shall return ±HUGE\_VAL, ±HUGE\_VALF, and ±HUGE\_VALL (having the same  
 22486 sign as the correct value), respectively.

22487 MX If  $x$  is NaN, a NaN shall be returned.22488 If  $x$  is 1 or 2, +0 shall be returned.22489 If  $x$  is ±Inf, +Inf shall be returned22490 **ERRORS**

22491 These functions shall fail if:

22492 Pole Error The  $x$  argument is a negative integer or zero.

22493 If the integer expression (math\_errhandling & MATH\_ERRNO) is non-zero, |  
 22494 then *errno* shall be set to [ERANGE]. If the integer expression |  
 22495 (math\_errhandling & MATH\_ERREXCEPT) is non-zero, then the divide-by- |  
 22496 zero floating-point exception shall be raised. |

22497 Range Error The result overflows

22498 If the integer expression (`math_errhandling & MATH_ERRNO`) is non-zero, |  
22499 then *errno* shall be set to [ERANGE]. If the integer expression |  
22500 (`math_errhandling & MATH_ERREXCEPT`) is non-zero, then the overflow |  
22501 floating-point exception shall be raised. |

**22502 EXAMPLES**

22503 None.

**22504 APPLICATION USAGE**

22505 On error, the expressions (`math_errhandling & MATH_ERRNO`) and (`math_errhandling &`  
22506 `MATH_ERREXCEPT`) are independent of each other, but at least one of them must be non-zero.

**22507 RATIONALE**

22508 None.

**22509 FUTURE DIRECTIONS**

22510 None.

**22511 SEE ALSO**

22512 *exp()*, *feclearexcept()*, *fetestexcept()*, *isnan()*, the Base Definitions volume of IEEE Std 1003.1-200x, |  
22513 Section 4.18, Treatment of Error Conditions for Mathematical Functions, <**math.h**> |

**22514 CHANGE HISTORY**

22515 First released in Issue 3.

**22516 Issue 5**

22517 The DESCRIPTION is updated to indicate how an application should check for an error. This  
22518 text was previously published in the APPLICATION USAGE section.

22519 A note indicating that this function need not be reentrant is added to the DESCRIPTION.

**22520 Issue 6**

22521 The *lgamma()* function is no longer marked as an extension.

22522 The *lgammaf()* and *lgammal()* functions are added for alignment with the ISO/IEC 9899:1999  
22523 standard.

22524 The DESCRIPTION, RETURN VALUE, ERRORS, and APPLICATION USAGE sections are  
22525 revised to align with the ISO/IEC 9899:1999 standard.

22526 IEC 60559:1989 standard floating-point extensions over the ISO/IEC 9899:1999 standard are  
22527 marked.

22528 XSI extensions are marked.

22529 **NAME**

22530 link — link to a file

22531 **SYNOPSIS**

22532 #include &lt;unistd.h&gt;

22533 int link(const char \*path1, const char \*path2);

22534 **DESCRIPTION**22535 The *link()* function shall create a new link (directory entry) for the existing file, *path1*.

22536 The *path1* argument points to a pathname naming an existing file. The *path2* argument points to  
 22537 a pathname naming the new directory entry to be created. The *link()* function shall atomically  
 22538 create a new link for the existing file and the link count of the file shall be incremented by one.

22539 If *path1* names a directory, *link()* shall fail unless the process has appropriate privileges and the  
 22540 implementation supports using *link()* on directories.

22541 Upon successful completion, *link()* shall mark for update the *st\_ctime* field of the file. Also, the  
 22542 *st\_ctime* and *st\_mtime* fields of the directory that contains the new entry shall be marked for  
 22543 update.

22544 If *link()* fails, no link shall be created and the link count of the file shall remain unchanged.

22545 The implementation may require that the calling process has permission to access the existing  
 22546 file.

22547 **RETURN VALUE**

22548 Upon successful completion, 0 shall be returned. Otherwise, -1 shall be returned and *errno* set to  
 22549 indicate the error.

22550 **ERRORS**22551 The *link()* function shall fail if:

22552 [EACCES] A component of either path prefix denies search permission, or the requested  
 22553 link requires writing in a directory that denies write permission, or the calling  
 22554 process does not have permission to access the existing file and this is  
 22555 required by the implementation.

22556 [EEXIST] The *path2* argument resolves to an existing file or refers to a symbolic link.

22557 [ELOOP] A loop exists in symbolic links encountered during resolution of the *path1* or  
 22558 *path2* argument.

22559 [EMLINK] The number of links to the file named by *path1* would exceed {LINK\_MAX}.

22560 [ENAMETOOLONG]

22561 The length of the *path1* or *path2* argument exceeds {PATH\_MAX} or a  
 22562 pathname component is longer than {NAME\_MAX}.

22563 [ENOENT] A component of either path prefix does not exist; the file named by *path1* does  
 22564 not exist; or *path1* or *path2* points to an empty string.

22565 [ENOSPC] The directory to contain the link cannot be extended.

22566 [ENOTDIR] A component of either path prefix is not a directory.

22567 [EPERM] The file named by *path1* is a directory and either the calling process does not  
 22568 have appropriate privileges or the implementation prohibits using *link()* on  
 22569 directories.

- 22570 [EROFS] The requested link requires writing in a directory on a read-only file system.
- 22571 [EXDEV] The link named by *path2* and the file named by *path1* are on different file  
22572 XSR systems and the implementation does not support links between file systems,  
22573 or *path1* refers to a named STREAM.
- 22574 The *link()* function may fail if:
- 22575 [ELOOP] More than {SYMLOOP\_MAX} symbolic links were encountered during  
22576 resolution of the *path1* or *path2* argument.
- 22577 [ENAMETOOLONG]  
22578 As a result of encountering a symbolic link in resolution of the *path1* or *path2* |  
22579 argument, the length of the substituted pathname string exceeded |  
22580 {PATH\_MAX}.

22581 **EXAMPLES**22582 **Creating a Link to a File**

22583 The following example shows how to create a link to a file named **/home/cnd/mod1** by creating a  
22584 new directory entry named **/modules/pass1**.

```
22585 #include <unistd.h>
22586
22586 char *path1 = "/home/cnd/mod1";
22587 char *path2 = "/modules/pass1";
22588 int status;
22589 ...
22590 status = link (path1, path2);
```

22591 **Creating a Link to a File Within a Program**

22592 In the following program example, the *link()* function links the **/etc/passwd** file (defined as |  
22593 **PASSWDFILE**) to a file named **/etc/opasswd** (defined as **SAVEFILE**), which is used to save the  
22594 current password file. Then, after removing the current password file (defined as  
22595 **PASSWDFILE**), the new password file is saved as the current password file using the *link()*  
22596 function again.

```
22597 #include <unistd.h>
22598
22598 #define LOCKFILE "/etc/ptmp"
22599 #define PASSWDFILE "/etc/passwd"
22600 #define SAVEFILE "/etc/opasswd"
22601 ...
22602 /* Save current password file */
22603 link (PASSWDFILE, SAVEFILE);
22604
22604 /* Remove current password file. */
22605 unlink (PASSWDFILE);
22606
22606 /* Save new password file as current password file. */
22607 link (LOCKFILE, PASSWDFILE);
```

22608 **APPLICATION USAGE**

22609 Some implementations do allow links between file systems.

**22610 RATIONALE**

22611 Linking to a directory is restricted to the superuser in most historical implementations because  
22612 this capability may produce loops in the file hierarchy or otherwise corrupt the file system. This  
22613 volume of IEEE Std 1003.1-200x continues that philosophy by prohibiting *link()* and *unlink()*  
22614 from doing this. Other functions could do it if the implementor designed such an extension.

22615 Some historical implementations allow linking of files on different file systems. Wording was  
22616 added to explicitly allow this optional behavior.

22617 The exception for cross-file system links is intended to apply only to links that are  
22618 programmatically indistinguishable from “hard” links.

**22619 FUTURE DIRECTIONS**

22620 None.

**22621 SEE ALSO**

22622 *symlink()*, *unlink()*, the Base Definitions volume of IEEE Std 1003.1-200x, <**unistd.h**>

**22623 CHANGE HISTORY**

22624 First released in Issue 1. Derived from Issue 1 of the SVID.

**22625 Issue 6**

22626 The following new requirements on POSIX implementations derive from alignment with the  
22627 Single UNIX Specification:

- 22628 • The [ELOOP] mandatory error condition is added.
- 22629 • A second [ENAMETOOLONG] is added as an optional error condition.

22630 The following changes were made to align with the IEEE P1003.1a draft standard:

- 22631 • An explanation is added of action when *path2* refers to a symbolic link.
- 22632 • The [ELOOP] optional error condition is added.

## 22633 NAME

22634 lio\_listio — list directed I/O (**REALTIME**)

## 22635 SYNOPSIS

22636 AIO #include &lt;aio.h&gt;

```
22637 int lio_listio(int mode, struct aiocb *restrict const list[restrict],
22638               int nent, struct sigevent *restrict sig);
22639
```

## 22640 DESCRIPTION

22641 The *lio\_listio()* function shall initiate a list of I/O requests with a single function call.

22642 The *mode* argument takes one of the values LIO\_WAIT or LIO\_NOWAIT declared in <aio.h> and  
 22643 determines whether the function returns when the I/O operations have been completed, or as  
 22644 soon as the operations have been queued. If the *mode* argument is LIO\_WAIT, the function shall  
 22645 wait until all I/O is complete and the *sig* argument shall be ignored.

22646 If the *mode* argument is LIO\_NOWAIT, the function shall return immediately, and asynchronous  
 22647 notification shall occur, according to the *sig* argument, when all the I/O operations complete. If  
 22648 *sig* is NULL, then no asynchronous notification shall occur. If *sig* is not NULL, asynchronous  
 22649 notification occurs as specified in Section 2.4.1 (on page 478) when all the requests in *list* have  
 22650 completed.

22651 The I/O requests enumerated by *list* are submitted in an unspecified order.

22652 The *list* argument is an array of pointers to **aiocb** structures. The array contains *nent* elements.  
 22653 The array may contain NULL elements, which shall be ignored.

22654 The *aio\_lio\_opcode* field of each **aiocb** structure specifies the operation to be performed. The  
 22655 supported operations are LIO\_READ, LIO\_WRITE, and LIO\_NOP; these symbols are defined in  
 22656 <aio.h>. The LIO\_NOP operation causes the list entry to be ignored. If the *aio\_lio\_opcode*  
 22657 element is equal to LIO\_READ, then an I/O operation is submitted as if by a call to *aio\_read()*  
 22658 with the *aiocbp* equal to the address of the **aiocb** structure. If the *aio\_lio\_opcode* element is equal  
 22659 to LIO\_WRITE, then an I/O operation is submitted as if by a call to *aio\_write()* with the *aiocbp*  
 22660 equal to the address of the **aiocb** structure.

22661 The *aio\_fildes* member specifies the file descriptor on which the operation is to be performed.22662 The *aio\_buf* member specifies the address of the buffer to or from which the data is transferred.22663 The *aio\_nbytes* member specifies the number of bytes of data to be transferred.

22664 The members of the **aiocb** structure further describe the I/O operation to be performed, in a  
 22665 manner identical to that of the corresponding **aiocb** structure when used by the *aio\_read()* and  
 22666 *aio\_write()* functions.

22667 The *nent* argument specifies how many elements are members of the list; that is, the length of the  
 22668 array.

22669 The behavior of this function is altered according to the definitions of synchronized I/O data  
 22670 integrity completion and synchronized I/O file integrity completion if synchronized I/O is  
 22671 enabled on the file associated with *aio\_fildes*.

22672 For regular files, no data transfer shall occur past the offset maximum established in the open  
 22673 file description associated with *aiocbp->aio\_fildes*.



## 22674 RETURN VALUE

22675 If the *mode* argument has the value LIO\_NOWAIT, the *lio\_listio()* function shall return the value  
 22676 zero if the I/O operations are successfully queued; otherwise, the function shall return the value  
 22677  $-1$  and set *errno* to indicate the error.

22678 If the *mode* argument has the value LIO\_WAIT, the *lio\_listio()* function shall return the value  
 22679 zero when all the indicated I/O has completed successfully. Otherwise, *lio\_listio()* shall return a  
 22680 value of  $-1$  and set *errno* to indicate the error.

22681 In either case, the return value only indicates the success or failure of the *lio\_listio()* call itself,  
 22682 not the status of the individual I/O requests. In some cases one or more of the I/O requests  
 22683 contained in the list may fail. Failure of an individual request does not prevent completion of  
 22684 any other individual request. To determine the outcome of each I/O request, the application  
 22685 shall examine the error status associated with each **aiocb** control block. The error statuses so  
 22686 returned are identical to those returned as the result of an *aio\_read()* or *aio\_write()* function.

## 22687 ERRORS

22688 The *lio\_listio()* function shall fail if:

22689 [EAGAIN] The resources necessary to queue all the I/O requests were not available. The  
 22690 application may check the error status for each **aiocb** to determine the  
 22691 individual request(s) that failed.

22692 [EAGAIN] The number of entries indicated by *nent* would cause the system-wide limit  
 22693 {AIO\_MAX} to be exceeded.

22694 [EINVAL] The *mode* argument is not a proper value, or the value of *nent* was greater than  
 22695 {AIO\_LISTIO\_MAX}.

22696 [EINTR] A signal was delivered while waiting for all I/O requests to complete during  
 22697 an LIO\_WAIT operation. Note that, since each I/O operation invoked by  
 22698 *lio\_listio()* may possibly provoke a signal when it completes, this error return  
 22699 may be caused by the completion of one (or more) of the very I/O operations  
 22700 being awaited. Outstanding I/O requests are not canceled, and the application  
 22701 shall examine each list element to determine whether the request was  
 22702 initiated, canceled, or completed.

22703 [EIO] One or more of the individual I/O operations failed. The application may  
 22704 check the error status for each **aiocb** structure to determine the individual  
 22705 request(s) that failed.

22706 In addition to the errors returned by the *lio\_listio()* function, if the *lio\_listio()* function succeeds  
 22707 or fails with errors of [EAGAIN], [EINTR], or [EIO], then some of the I/O specified by the list  
 22708 may have been initiated. If the *lio\_listio()* function fails with an error code other than [EAGAIN],  
 22709 [EINTR], or [EIO], no operations from the list shall have been initiated. The I/O operation  
 22710 indicated by each list element can encounter errors specific to the individual read or write  
 22711 function being performed. In this event, the error status for each **aiocb** control block contains the  
 22712 associated error code. The error codes that can be set are the same as would be set by a *read()* or  
 22713 *write()* function, with the following additional error codes possible:

22714 [EAGAIN] The requested I/O operation was not queued due to resource limitations.

22715 [ECANCELED] The requested I/O was canceled before the I/O completed due to an explicit  
 22716 *aio\_cancel()* request.

22717 [EFBIG] The *aiocbp->aio\_lio\_opcode* is LIO\_WRITE, the file is a regular file, *aiocbp->aio\_nbytes*  
 22718 is greater than 0, and the *aiocbp->aio\_offset* is greater than or equal  
 22719 to the offset maximum in the open file description associated with *aiocbp-*

22720 *> aio\_fildes.*

22721 [EINPROGRESS] The requested I/O is in progress.

22722 [EOVERFLOW] The *aiocbp->aio\_lio\_opcode* is LIO\_READ, the file is a regular file, *aiocbp->aio\_nbytes* is greater than 0, and the *aiocbp->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 *aiocbp->aio\_fildes*.

#### 22726 EXAMPLES

22727 None.

#### 22728 APPLICATION USAGE

22729 None.

#### 22730 RATIONALE

22731 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]).

22737 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.

#### 22740 FUTURE DIRECTIONS

22741 None.

#### 22742 SEE ALSO

22743 *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**>

#### 22745 CHANGE HISTORY

22746 First released in Issue 5. Included for alignment with the POSIX Realtime Extension.

22747 Large File Summit extensions are added.

#### 22748 Issue 6

22749 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.

22751 The *lio\_listio()* function is marked as part of the Asynchronous Input and Output option.

22752 The following new requirements on POSIX implementations derive from alignment with the Single UNIX Specification:

- 22754 • 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 *aiocbp->aio\_fildes*. This change is to support large files.
- 22755 • The [EBIG] and [EOVERFLOW] error conditions are defined. This change is to support large files.

22759 The DESCRIPTION is updated to avoid use of the term “must” for application requirements.

22760 The **restrict** keyword is added to the *lio\_listio()* prototype for alignment with the ISO/IEC 9899:1999 standard.

22761

22762 **NAME**

22763 listen — listen for socket connections and limit the queue of incoming connections

22764 **SYNOPSIS**

```
22765 #include <sys/socket.h>
22766 int listen(int socket, int backlog);
```

22767 **DESCRIPTION**

22768 The *listen()* function shall mark a connection-mode socket, specified by the *socket* argument, as  
 22769 accepting connections.

22770 The *backlog* argument provides a hint to the implementation which the implementation shall use  
 22771 to limit the number of outstanding connections in the socket's listen queue. Implementations  
 22772 may impose a limit on *backlog* and silently reduce the specified value. Normally, a larger *backlog*  
 22773 argument value shall result in a larger or equal length of the listen queue. Implementations shall  
 22774 support values of *backlog* up to SOMAXCONN, defined in <sys/socket.h>.

22775 The implementation may include incomplete connections in its listen queue. The limits on the  
 22776 number of incomplete connections and completed connections queued may be different.

22777 The implementation may have an upper limit on the length of the listen queue—either global or  
 22778 per accepting socket. If *backlog* exceeds this limit, the length of the listen queue is set to the limit.

22779 If *listen()* is called with a *backlog* argument value that is less than 0, the function behaves as if it  
 22780 had been called with a *backlog* argument value of 0.

22781 A *backlog* argument of 0 may allow the socket to accept connections, in which case the length of  
 22782 the listen queue may be set to an implementation-defined minimum value.

22783 The socket in use may require the process to have appropriate privileges to use the *listen()*  
 22784 function.

22785 **RETURN VALUE**

22786 Upon successful completions, *listen()* shall return 0; otherwise, -1 shall be returned and *errno* set  
 22787 to indicate the error.

22788 **ERRORS**

22789 The *listen()* function shall fail if:

22790 [EBADF] The *socket* argument is not a valid file descriptor.

22791 [EDESTADDRREQ]

22792 The socket is not bound to a local address, and the protocol does not support  
 22793 listening on an unbound socket.

22794 [EINVAL] The *socket* is already connected.

22795 [ENOTSOCK] The *socket* argument does not refer to a socket.

22796 [EOPNOTSUPP] The socket protocol does not support *listen()*.

22797 The *listen()* function may fail if:

22798 [EACCES] The calling process does not have the appropriate privileges.

22799 [EINVAL] The *socket* has been shut down.

22800 [ENOBUFS] Insufficient resources are available in the system to complete the call.

22801 **EXAMPLES**

22802           None.

22803 **APPLICATION USAGE**

22804           None.

22805 **RATIONALE**

22806           None.

22807 **FUTURE DIRECTIONS**

22808           None.

22809 **SEE ALSO**

22810           *accept()*, *connect()*, *socket()*, the Base Definitions volume of IEEE Std 1003.1-200x, <**sys/socket.h**>

22811 **CHANGE HISTORY**

22812           First released in Issue 6. Derived from the XNS, Issue 5.2 specification.

22813           The DESCRIPTION is updated to describe the relationship of SOMAXCONN and the *backlog*

22814           argument.

22815 **NAME**

22816       llabs — return a long integer absolute value

22817 **SYNOPSIS**

22818       #include <stdlib.h>

22819       long long llabs(long long *i*);

22820 **DESCRIPTION**

22821       Refer to *labs()*.

22822 **NAME**

22823        `lldiv` — compute quotient and remainder of a long division

22824 **SYNOPSIS**

22825        `#include <stdlib.h>`

22826        `lldiv_t lldiv(long long numer, long long denom);`

22827 **DESCRIPTION**

22828        Refer to *ldiv()*.

22829 **NAME**

22830 llrint, llrintf, llrintl, — round to nearest integer value using current rounding direction

22831 **SYNOPSIS**

22832 #include <math.h>

22833 long long llrint(double x);

22834 long long llrintf(float x);

22835 long long llrintl(long double x);

22836 **DESCRIPTION**

22837 **CX** The functionality described on this reference page is aligned with the ISO C standard. Any  
 22838 conflict between the requirements described here and the ISO C standard is unintentional. This  
 22839 volume of IEEE Std 1003.1-200x defers to the ISO C standard.

22840 These functions shall round their argument to the nearest integer value, rounding according to  
 22841 the current rounding direction.

22842 An application wishing to check for error situations should set *errno* to zero and call  
 22843 *feclearexcept*(FE\_ALL\_EXCEPT) before calling these functions. On return, if *errno* is non-zero or  
 22844 *fetetestexcept*(FE\_INVALID | FE\_DIVBYZERO | FE\_OVERFLOW | FE\_UNDERFLOW) is non-  
 22845 zero, an error has occurred.

22846 **RETURN VALUE**

22847 Upon successful completion, these functions shall return the rounded integer value.

22848 **MX** If *x* is NaN, a domain error shall occur, and an unspecified value is returned.

22849 If *x* is +Inf, a domain error shall occur and an unspecified value is returned.

22850 If *x* is -Inf, a domain error shall occur and an unspecified value is returned.

22851 If the correct value is positive and too large to represent as a **long long**, a domain error shall  
 22852 occur and an unspecified value is returned.

22853 If the correct value is negative and too large to represent as a **long long**, a domain error shall  
 22854 occur and an unspecified value is returned.

22855 **ERRORS**

22856 These functions shall fail if:

22857 **MX** Domain Error The *x* argument is NaN or ±Inf, or the correct value is not representable as an  
 22858 integer.

22859 If the integer expression (math\_errhandling & MATH\_ERRNO) is non-zero,  
 22860 then *errno* shall be set to [EDOM]. If the integer expression (math\_errhandling  
 22861 & MATH\_ERREXCEPT) is non-zero, then the invalid floating-point exception  
 22862 shall be raised.

22863 **EXAMPLES**

22864 None.

22865 **APPLICATION USAGE**

22866 On error, the expressions (math\_errhandling & MATH\_ERRNO) and (math\_errhandling &  
 22867 MATH\_ERREXCEPT) are independent of each other, but at least one of them must be non-zero.

22868 **RATIONALE**

22869 These functions provide floating-to-integer conversions. They round according to the current  
 22870 rounding direction. If the rounded value is outside the range of the return type, the numeric  
 22871 result is unspecified and the invalid floating-point exception is raised. When they raise no other  
 22872 floating-point exception and the result differs from the argument, they raise the inexact

22873 floating-point exception.

22874 **FUTURE DIRECTIONS**

22875 None.

22876 **SEE ALSO**

22877 *feclearexcept()*, *fetestexcept()*, *lrint()*, the Base Definitions volume of IEEE Std 1003.1-200x, Section |  
22878 4.18, Treatment of Error Conditions for Mathematical Functions, <math.h> |

22879 **CHANGE HISTORY**

22880 First released in Issue 6. Derived from the ISO/IEC 9899:1999 standard.



22881 **NAME**

22882 llround, llroundf, llroundl, — round to nearest integer value

22883 **SYNOPSIS**

22884 #include &lt;math.h&gt;

22885 long long llround(double x);

22886 long long llroundf(float x);

22887 long long llroundl(long double x);

22888 **DESCRIPTION**

22889 **CX** The functionality described on this reference page is aligned with the ISO C standard. Any  
 22890 conflict between the requirements described here and the ISO C standard is unintentional. This  
 22891 volume of IEEE Std 1003.1-200x defers to the ISO C standard.

22892 These functions shall round their argument to the nearest integer value, rounding halfway cases  
 22893 away from zero, regardless of the current rounding direction.

22894 An application wishing to check for error situations should set *errno* to zero and call  
 22895 *feclearexcept*(FE\_ALL\_EXCEPT) before calling these functions. On return, if *errno* is non-zero or  
 22896 *fetestexcept*(FE\_INVALID | FE\_DIVBYZERO | FE\_OVERFLOW | FE\_UNDERFLOW) is non-  
 22897 zero, an error has occurred.

22898 **RETURN VALUE**

22899 Upon successful completion, these functions shall return the rounded integer value.

22900 **MX** If *x* is NaN, a domain error shall occur, and an unspecified value is returned.22901 If *x* is +Inf, a domain error shall occur and an unspecified value is returned.22902 If *x* is -Inf, a domain error shall occur and an unspecified value is returned.

22903 If the correct value is positive and too large to represent as a **long long**, a domain error shall  
 22904 occur and an unspecified value is returned.

22905 If the correct value is negative and too large to represent as a **long long**, a domain error shall  
 22906 occur and an unspecified value is returned.

22907 **ERRORS**

22908 These functions shall fail if:

22909 **MX** Domain Error The *x* argument is NaN or ±Inf, or the correct value is not representable as an  
 22910 integer.

22911 If the integer expression (math\_errhandling & MATH\_ERRNO) is non-zero,  
 22912 then *errno* shall be set to [EDOM]. If the integer expression (math\_errhandling  
 22913 & MATH\_ERREXCEPT) is non-zero, then the invalid floating-point exception  
 22914 shall be raised.

22915 **EXAMPLES**

22916 None.

22917 **APPLICATION USAGE**

22918 On error, the expressions (math\_errhandling & MATH\_ERRNO) and (math\_errhandling &  
 22919 MATH\_ERREXCEPT) are independent of each other, but at least one of them must be non-zero.

22920 **RATIONALE**

22921 These functions provide floating-to-integer conversions. They round according to the current  
 22922 rounding direction. If the rounded value is outside the range of the return type, the numeric  
 22923 result is unspecified and the invalid floating-point exception is raised. When they raise no other  
 22924 floating-point exception and the result differs from the argument, they raise the inexact

22925 floating-point exception.

22926 These functions differ from the *llrint()* functions in that the default rounding direction for the  
22927 *lround()* functions round halfway cases away from zero and need not raise the inexact floating-  
22928 point exception for non-integer arguments that round to within the range of the return type.

22929 **FUTURE DIRECTIONS**

22930 None.

22931 **SEE ALSO**

22932 *feclearexcept()*, *fetetestexcept()*, *lround()*, the Base Definitions volume of IEEE Std 1003.1-200x, |  
22933 Section 4.18, Treatment of Error Conditions for Mathematical Functions, <math.h> |

22934 **CHANGE HISTORY**

22935 First released in Issue 6. Derived from the ISO/IEC 9899:1999 standard.

22936 **NAME**

22937 localeconv — return locale-specific information

22938 **SYNOPSIS**

22939 #include &lt;locale.h&gt;

22940 struct lconv \*localeconv(void);

22941 **DESCRIPTION**

22942 **CX** The functionality described on this reference page is aligned with the ISO C standard. Any  
 22943 conflict between the requirements described here and the ISO C standard is unintentional. This  
 22944 volume of IEEE Std 1003.1-200x defers to the ISO C standard.

22945 The *localeconv()* function shall set the components of an object with the type **struct lconv** with  
 22946 the values appropriate for the formatting of numeric quantities (monetary and otherwise)  
 22947 according to the rules of the current locale.

22948 The members of the structure with type **char \*** are pointers to strings, any of which (except  
 22949 **decimal\_point**) can point to " ", to indicate that the value is not available in the current locale or  
 22950 is of zero length. The members with type **char** are non-negative numbers, any of which can be  
 22951 {CHAR\_MAX} to indicate that the value is not available in the current locale.

22952 The members include the following:

22953 **char \*decimal\_point**

22954 The radix character used to format non-monetary quantities.

22955 **char \*thousands\_sep**

22956 The character used to separate groups of digits before the decimal-point character in  
 22957 formatted non-monetary quantities.

22958 **char \*grouping**

22959 A string whose elements taken as one-byte integer values indicate the size of each group of  
 22960 digits in formatted non-monetary quantities.

22961 **char \*int\_curr\_symbol**

22962 The international currency symbol applicable to the current locale. The first three  
 22963 characters contain the alphabetic international currency symbol in accordance with those  
 22964 specified in the ISO 4217:1995 standard. The fourth character (immediately preceding the  
 22965 null byte) is the character used to separate the international currency symbol from the  
 22966 monetary quantity.

22967 **char \*currency\_symbol**

22968 The local currency symbol applicable to the current locale.

22969 **char \*mon\_decimal\_point**

22970 The radix character used to format monetary quantities.

22971 **char \*mon\_thousands\_sep**

22972 The separator for groups of digits before the decimal-point in formatted monetary  
 22973 quantities.

22974 **char \*mon\_grouping**

22975 A string whose elements taken as one-byte integer values indicate the size of each group of  
 22976 digits in formatted monetary quantities.

22977 **char \*positive\_sign**

22978 The string used to indicate a non-negative valued formatted monetary quantity.

- 22979 **char \*negative\_sign**  
 22980 The string used to indicate a negative valued formatted monetary quantity.
- 22981 **char int\_frac\_digits**  
 22982 The number of fractional digits (those after the decimal-point) to be displayed in an  
 22983 internationally formatted monetary quantity.
- 22984 **char frac\_digits**  
 22985 The number of fractional digits (those after the decimal-point) to be displayed in a  
 22986 formatted monetary quantity.
- 22987 **char p\_cs\_precedes**  
 22988 Set to 1 if the **currency\_symbol** or **int\_curr\_symbol** precedes the value for a non-negative  
 22989 formatted monetary quantity. Set to 0 if the symbol succeeds the value.
- 22990 **char p\_sep\_by\_space**  
 22991 Set to 0 if no space separates the **currency\_symbol** or **int\_curr\_symbol** from the value for a  
 22992 non-negative formatted monetary quantity. Set to 1 if a space separates the symbol from the  
 22993 value; and set to 2 if a space separates the symbol and the sign string, if adjacent. XSI
- 22994 **char n\_cs\_precedes**  
 22995 Set to 1 if the **currency\_symbol** or **int\_curr\_symbol** precedes the value for a negative  
 22996 formatted monetary quantity. Set to 0 if the symbol succeeds the value.
- 22997 **char n\_sep\_by\_space**  
 22998 Set to 0 if no space separates the **currency\_symbol** or **int\_curr\_symbol** from the value for a  
 22999 negative formatted monetary quantity. Set to 1 if a space separates the symbol from the  
 23000 value; and set to 2 if a space separates the symbol and the sign string, if adjacent. XSI
- 23001 **char p\_sign\_posn**  
 23002 Set to a value indicating the positioning of the **positive\_sign** for a non-negative formatted  
 23003 monetary quantity.
- 23004 **char n\_sign\_posn**  
 23005 Set to a value indicating the positioning of the **negative\_sign** for a negative formatted  
 23006 monetary quantity.
- 23007 **char int\_p\_cs\_precedes**  
 23008 Set to 1 or 0 if the **int\_curr\_symbol** respectively precedes or succeeds the value for a non- |  
 23009 negative internationally formatted monetary quantity.
- 23010 **char int\_n\_cs\_precedes**  
 23011 Set to 1 or 0 if the **int\_curr\_symbol** respectively precedes or succeeds the value for a |  
 23012 negative internationally formatted monetary quantity.
- 23013 **char int\_p\_sep\_by\_space**  
 23014 Set to a value indicating the separation of the **int\_curr\_symbol**, the sign string, and the |  
 23015 value for a non-negative internationally formatted monetary quantity.
- 23016 **char int\_n\_sep\_by\_space**  
 23017 Set to a value indicating the separation of the **int\_curr\_symbol**, the sign string, and the |  
 23018 value for a negative internationally formatted monetary quantity.
- 23019 **char int\_p\_sign\_posn**  
 23020 Set to a value indicating the positioning of the **positive\_sign** for a non-negative  
 23021 internationally formatted monetary quantity.
- 23022 **char int\_n\_sign\_posn**  
 23023 Set to a value indicating the positioning of the **negative\_sign** for a negative internationally

- 23024 formatted monetary quantity.
- 23025 The elements of **grouping** and **mon\_grouping** are interpreted according to the following:
- 23026 {CHAR\_MAX} No further grouping is to be performed.
- 23027 0 The previous element is to be repeatedly used for the remainder of the digits.
- 23028 *other* The integer value is the number of digits that comprise the current group. The  
23029 next element is examined to determine the size of the next group of digits  
23030 before the current group.
- 23031 The values of **p\_sep\_by\_space**, **n\_sep\_by\_space**, **int\_p\_sep\_by\_space**, and **int\_n\_sep\_by\_space**  
23032 are interpreted according to the following:
- 23033 0 No space separates the currency symbol and value.
- 23034 1 If the currency symbol and sign string are adjacent, a space separates them from the value;  
23035 otherwise, a space separates the currency symbol from the value.
- 23036 2 If the currency symbol and sign string are adjacent, a space separates them; otherwise, a  
23037 space separates the sign string from the value.
- 23038 For **int\_p\_sep\_by\_space** and **int\_n\_sep\_by\_space**, the fourth character of **int\_curr\_symbol** is |  
23039 used instead of a space. |
- 23040 The values of **p\_sign\_posn**, **n\_sign\_posn**, **int\_p\_sign\_posn**, and **int\_n\_sign\_posn** are |  
23041 interpreted according to the following:
- 23042 0 Parentheses surround the quantity and **currency\_symbol** or **int\_curr\_symbol**.
- 23043 1 The sign string precedes the quantity and **currency\_symbol** or **int\_curr\_symbol**.
- 23044 2 The sign string succeeds the quantity and **currency\_symbol** or **int\_curr\_symbol**.
- 23045 3 The sign string immediately precedes the **currency\_symbol** or **int\_curr\_symbol**.
- 23046 4 The sign string immediately succeeds the **currency\_symbol** or **int\_curr\_symbol**.
- 23047 The implementation shall behave as if no function in this volume of IEEE Std 1003.1-200x calls  
23048 *localeconv()*.
- 23049 cx The *localeconv()* function need not be reentrant. A function that is not required to be reentrant is  
23050 not required to be thread-safe.
- 23051 **RETURN VALUE**
- 23052 The *localeconv()* function shall return a pointer to the filled-in object. The application shall not  
23053 modify the structure pointed to by the return value which may be overwritten by a subsequent  
23054 call to *localeconv()*. In addition, calls to *setlocale()* with the categories *LC\_ALL*, *LC\_MONETARY*,  
23055 or *LC\_NUMERIC* may overwrite the contents of the structure.
- 23056 **ERRORS**
- 23057 No errors are defined.

23058 **EXAMPLES**

23059 None.

23060 **APPLICATION USAGE**

23061 The following table illustrates the rules which may be used by four countries to format monetary  
 23062 quantities.

23063

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	SFrS.1,234.56	SFrS.1,234.56C	CHF 1,234.56

23064

23065

23066

23067

23068 For these four countries, the respective values for the monetary members of the structure  
 23069 returned by *localeconv()* are:

23070

	Italy	Netherlands	Norway	Switzerland
23071 <b>int_curr_symbol</b>	"ITL. "	"NLG "	"NOK "	"CHF "
23072 <b>currency_symbol</b>	"L. "	"F "	"kr "	"SFrS. "
23073 <b>mon_decimal_point</b>	" "	","	","	."
23074 <b>mon_thousands_sep</b>	"."	"."	"."	","
23075 <b>mon_grouping</b>	"\3 "	"\3 "	"\3 "	"\3 "
23076 <b>positive_sign</b>	" "	" "	" "	" "
23077 <b>negative_sign</b>	"- "	"- "	"- "	"C "
23078 <b>int_frac_digits</b>	0	2	2	2
23079 <b>frac_digits</b>	0	2	2	2
23080 <b>p_cs_precedes</b>	1	1	1	1
23081 <b>p_sep_by_space</b>	0	1	0	0
23082 <b>n_cs_precedes</b>	1	1	1	1
23083 <b>n_sep_by_space</b>	0	1	0	0
23084 <b>p_sign_posn</b>	1	1	1	1
23085 <b>n_sign_posn</b>	1	4	2	2
23086 <b>int_p_cs_precedes</b>	1	1	1	1
23087 <b>int_n_cs_precedes</b>	1	1	1	1
23088 <b>int_p_sep_by_space</b>	0	0	0	0
23089 <b>int_n_sep_by_space</b>	0	0	0	0
23090 <b>int_p_sign_posn</b>	1	1	1	1
23091 <b>int_n_sign_posn</b>	1	4	4	2

23092 **RATIONALE**

23093 None.

23094 **FUTURE DIRECTIONS**

23095 None.

23096 **SEE ALSO**

23097 *isalpha()*, *isascii()*, *nl\_langinfo()*, *printf()*, *scanf()*, *setlocale()*, *strcat()*, *strchr()*, *strcmp()*, *strcoll()*,  
 23098 *strcpy()*, *strftime()*, *strlen()*, *strpbrk()*, *strspn()*, *strtok()*, *strxfrm()*, *strtod()*, the Base Definitions  
 23099 volume of IEEE Std 1003.1-200x, <langinfo.h>, <locale.h>

23100 **CHANGE HISTORY**

23101 First released in Issue 4. Derived from the ANSI C standard.

23102 **Issue 6**

23103 A note indicating that this function need not be reentrant is added to the DESCRIPTION.

23104 The RETURN VALUE section is rewritten to avoid use of the term “must”.

23105 This reference page is updated for alignment with the ISO/IEC 9899: 1999 standard. |

23106 ISO/IEC 9899: 1999 standard, Technical Corrigendum No. 1 is incorporated. |

## 23107 NAME

23108 localtime, localtime\_r — convert a time value to a broken-down local time

## 23109 SYNOPSIS

23110 #include <time.h>

23111 struct tm \*localtime(const time\_t \*timer);

23112 TSF struct tm \*localtime\_r(const time\_t \*restrict timer,

23113 struct tm \*restrict result);

23114

## 23115 DESCRIPTION

23116 CX For *localtime()*: The functionality described on this reference page is aligned with the ISO C  
23117 standard. Any conflict between the requirements described here and the ISO C standard is  
23118 unintentional. This volume of IEEE Std 1003.1-200x defers to the ISO C standard.

23119 The *localtime()* function shall convert the time in seconds since the Epoch pointed to by *timer*  
23120 into a broken-down time, expressed as a local time. The function corrects for the timezone and  
23121 CX any seasonal time adjustments. Local timezone information is used as though *localtime()* calls  
23122 *tzset()*.

23123 The relationship between a time in seconds since the Epoch used as an argument to *localtime()* |  
23124 and the **tm** structure (defined in the <time.h> header) is that the result shall be as specified in the |  
23125 expression given in the definition of seconds since the Epoch (see the Base Definitions volume of |  
23126 IEEE Std 1003.1-200x, Section 4.14, Seconds Since the Epoch) corrected for timezone and any |  
23127 seasonal time adjustments, where the names in the structure and in the expression correspond. |

23128 TSF The same relationship shall apply for *localtime\_r()*. |

23129 CX The *localtime()* function need not be reentrant. A function that is not required to be reentrant is |  
23130 not required to be thread-safe. |

23131 The *asctime()*, *ctime()*, *gmtime()*, and *localtime()* functions shall return values in one of two static |  
23132 objects: a broken-down time structure and an array of type **char**. Execution of any of the |  
23133 functions may overwrite the information returned in either of these objects by any of the other |  
23134 functions. |

23135 TSF The *localtime\_r()* function shall convert the time in seconds since the Epoch pointed to by *timer*  
23136 into a broken-down time stored in the structure to which *result* points. The *localtime\_r()* function  
23137 shall also return a pointer to that same structure.

23138 Unlike *localtime()*, the reentrant version is not required to set *tzname*.

## 23139 RETURN VALUE

23140 The *localtime()* function shall return a pointer to the broken-down time structure.

23141 TSF Upon successful completion, *localtime\_r()* shall return a pointer to the structure pointed to by  
23142 the argument *result*.

## 23143 ERRORS

23144 No errors are defined.



23145 **EXAMPLES**23146 **Getting the Local Date and Time**

23147 The following example uses the *time()* function to calculate the time elapsed, in seconds, since  
 23148 January 1, 1970 0:00 UTC (the Epoch), *localtime()* to convert that value to a broken-down time,  
 23149 and *asctime()* to convert the broken-down time values into a printable string.

```
23150 #include <stdio.h>
23151 #include <time.h>
23152 main()
23153 {
23154     time_t result;
23155     result = time(NULL);
23156     printf("%s%ld secs since the Epoch\n",
23157           asctime(localtime(&result)),
23158           (long)result);
23159     return(0);
23160 }
```

23161 This example writes the current time to *stdout* in a form like this:

```
23162 Wed Jun 26 10:32:15 1996
23163 835810335 secs since the Epoch
```

23164 **Getting the Modification Time for a File**

23165 The following example gets the modification time for a file. The *localtime()* function converts the  
 23166 **time\_t** value of the last modification date, obtained by a previous call to *stat()*, into a **tm**  
 23167 structure that contains the year, month, day, and so on.

```
23168 #include <time.h>
23169 ...
23170 struct stat statbuf;
23171 ...
23172 tm = localtime(&statbuf.st_mtime);
23173 ...
```

23174 **Timing an Event**

23175 The following example gets the current time, converts it to a string using *localtime()* and  
 23176 *asctime()*, and prints it to standard output using *fputs()*. It then prints the number of minutes to  
 23177 an event being timed.

```
23178 #include <time.h>
23179 #include <stdio.h>
23180 ...
23181 time_t now;
23182 int minutes_to_event;
23183 ...
23184 time(&now);
23185 printf("The time is ");
23186 fputs(asctime(localtime(&now)), stdout);
23187 printf("There are still %d minutes to the event.\n",
```

23188                   minutes\_to\_event);

23189                   ...

23190 **APPLICATION USAGE**

23191                   The *localtime\_r()* function is thread-safe and returns values in a user-supplied buffer instead of  
23192                   possibly using a static data area that may be overwritten by each call.

23193 **RATIONALE**

23194                   None.

23195 **FUTURE DIRECTIONS**

23196                   None.

23197 **SEE ALSO**

23198                   *asctime()*, *clock()*, *ctime()*, *difftime()*, *getdate()*, *gmtime()*, *mktime()*, *strftime()*, *strptime()*, *time()*,  
23199                   *utime()*, the Base Definitions volume of IEEE Std 1003.1-200x, <**time.h**>

23200 **CHANGE HISTORY**

23201                   First released in Issue 1. Derived from Issue 1 of the SVID.

23202 **Issue 5**

23203                   A note indicating that the *localtime()* function need not be reentrant is added to the  
23204                   DESCRIPTION.

23205                   The *localtime\_r()* function is included for alignment with the POSIX Threads Extension.

23206 **Issue 6**

23207                   The *localtime\_r()* function is marked as part of the Thread-Safe Functions option.

23208                   Extensions beyond the ISO C standard are now marked.

23209                   The APPLICATION USAGE section is updated to include a note on the thread-safe function and  
23210                   its avoidance of possibly using a static data area.

23211                   The **restrict** keyword is added to the *localtime\_r()* prototype for alignment with the  
23212                   ISO/IEC 9899:1999 standard.

23213                   Examples are added.

23214 **NAME**

23215 lockf — record locking on files

23216 **SYNOPSIS**

23217 XSI #include &lt;unistd.h&gt;

23218 int lockf(int *fildes*, int *function*, off\_t *size*);

23219

23220 **DESCRIPTION**

23221 The *lockf()* function shall lock sections of a file with advisory-mode locks. Calls to *lockf()* from  
 23222 other threads which attempt to lock the locked file section shall either return an error value or  
 23223 block until the section becomes unlocked. All the locks for a process are removed when the  
 23224 process terminates. Record locking with *lockf()* shall be supported for regular files and may be  
 23225 supported for other files.

23226 The *fildes* argument is an open file descriptor. To establish a lock with this function, the file  
 23227 descriptor shall be opened with write-only permission (O\_WRONLY) or with read/write  
 23228 permission (O\_RDWR).

23229 The *function* argument is a control value which specifies the action to be taken. The permissible  
 23230 values for *function* are defined in <unistd.h> as follows:

23231

23232

23233

23234

23235

23236

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.

23237 F\_TEST shall detect if a lock by another process is present on the specified section.

23238 F\_LOCK and F\_TLOCK shall both lock a section of a file if the section is available.

23239 F\_ULOCK shall remove locks from a section of the file.

23240 The *size* argument is the number of contiguous bytes to be locked or unlocked. The section to be  
 23241 locked or unlocked starts at the current offset in the file and extends forward for a positive *size*  
 23242 or backward for a negative *size* (the preceding bytes up to but not including the current offset).  
 23243 If *size* is 0, the section from the current offset through the largest possible file offset shall be  
 23244 locked (that is, from the current offset through the present or any future end-of-file). An area  
 23245 need not be allocated to the file to be locked because locks may exist past the end-of-file.

23246 The sections locked with F\_LOCK or F\_TLOCK may, in whole or in part, contain or be contained  
 23247 by a previously locked section for the same process. When this occurs, or if adjacent locked  
 23248 sections would occur, the sections shall be combined into a single locked section. If the request  
 23249 would cause the number of locks to exceed a system-imposed limit, the request shall fail.

23250 F\_LOCK and F\_TLOCK requests differ only by the action taken if the section is not available.  
 23251 F\_LOCK shall block the calling thread until the section is available. F\_TLOCK shall cause the  
 23252 function to fail if the section is already locked by another process.

23253 File locks shall be released on first close by the locking process of any file descriptor for the file.

23254 F\_ULOCK requests may release (wholly or in part) one or more locked sections controlled by the  
 23255 process. Locked sections shall be unlocked starting at the current file offset through *size* bytes or  
 23256 to the end-of-file if *size* is (off\_t)0. When all of a locked section is not released (that is, when the  
 23257 beginning or end of the area to be unlocked falls within a locked section), the remaining portions  
 23258 of that section shall remain locked by the process. Releasing the center portion of a locked

- 23259 section shall cause the remaining locked beginning and end portions to become two separate  
 23260 locked sections. If the request would cause the number of locks in the system to exceed a  
 23261 system-imposed limit, the request shall fail.
- 23262 A potential for deadlock occurs if the threads of a process controlling a locked section are  
 23263 blocked by accessing another process' locked section. If the system detects that deadlock would  
 23264 occur, *lockf()* shall fail with an [EDEADLK] error.
- 23265 The interaction between *fcntl()* and *lockf()* locks is unspecified.
- 23266 Blocking on a section shall be interrupted by any signal.
- 23267 An F\_ULOCK request in which *size* is non-zero and the offset of the last byte of the requested  
 23268 section is the maximum value for an object of type *off\_t*, when the process has an existing lock  
 23269 in which *size* is 0 and which includes the last byte of the requested section, shall be treated as a  
 23270 request to unlock from the start of the requested section with a size equal to 0. Otherwise, an  
 23271 F\_ULOCK request shall attempt to unlock only the requested section.
- 23272 Attempting to lock a section of a file that is associated with a buffered stream produces  
 23273 unspecified results.
- 23274 **RETURN VALUE**
- 23275 Upon successful completion, *lockf()* shall return 0. Otherwise, it shall return -1, set *errno* to  
 23276 indicate an error, and existing locks shall not be changed.
- 23277 **ERRORS**
- 23278 The *lockf()* function shall fail if:
- 23279 [EBADF] The *fildev* argument is not a valid open file descriptor; or *function* is F\_LOCK  
 23280 or F\_TLOCK and *fildev* is not a valid file descriptor open for writing.
- 23281 [EACCES] or [EAGAIN]  
 23282 The *function* argument is F\_TLOCK or F\_TEST and the section is already  
 23283 locked by another process.
- 23284 [EDEADLK] The *function* argument is F\_LOCK and a deadlock is detected.
- 23285 [EINTR] A signal was caught during execution of the function.
- 23286 [EINVAL] The *function* argument is not one of F\_LOCK, F\_TLOCK, F\_TEST, or  
 23287 F\_ULOCK; or *size* plus the current file offset is less than 0.
- 23288 [EOVERFLOW] The offset of the first, or if *size* is not 0 then the last, byte in the requested  
 23289 section cannot be represented correctly in an object of type *off\_t*.
- 23290 The *lockf()* function may fail if:
- 23291 [EAGAIN] The *function* argument is F\_LOCK or F\_TLOCK and the file is mapped with  
 23292 *mmap()*.
- 23293 [EDEADLK] or [ENOLCK]  
 23294 The *function* argument is F\_LOCK, F\_TLOCK, or F\_ULOCK, and the request  
 23295 would cause the number of locks to exceed a system-imposed limit.
- 23296 [EOPNOTSUPP] or [EINVAL]  
 23297 The implementation does not support the locking of files of the type indicated  
 23298 by the *fildev* argument.

23299 **EXAMPLES**23300 **Locking a Portion of a File**

23301 In the following example, a file named `/home/cnd/mod1` is being modified. Other processes that  
 23302 use locking are prevented from changing it during this process. Only the first 10,000 bytes are  
 23303 locked, and the lock call fails if another process has any part of this area locked already.

```
23304 #include <fcntl.h>
23305 #include <unistd.h>
23306 int fildes;
23307 int status;
23308 ...
23309 fildes = open("/home/cnd/mod1", O_RDWR);
23310 status = lockf(fildes, F_TLOCK, (off_t)10000);
```

23311 **APPLICATION USAGE**

23312 Record-locking should not be used in combination with the *fopen()*, *fread()*, *fwrite()*, and other  
 23313 *stdio* functions. Instead, the more primitive, non-buffered functions (such as *open()*) should be  
 23314 used. Unexpected results may occur in processes that do buffering in the user address space. The  
 23315 process may later read/write data which is/was locked. The *stdio* functions are the most  
 23316 common source of unexpected buffering.

23317 The *alarm()* function may be used to provide a timeout facility in applications requiring it.

23318 **RATIONALE**

23319 None.

23320 **FUTURE DIRECTIONS**

23321 None.

23322 **SEE ALSO**

23323 *alarm()*, *chmod()*, *close()*, *creat()*, *fcntl()*, *fopen()*, *mmap()*, *open()*, *read()*, *write()*, the Base  
 23324 Definitions volume of IEEE Std 1003.1-200x, **<unistd.h>**

23325 **CHANGE HISTORY**

23326 First released in Issue 4, Version 2.

23327 **Issue 5**

23328 Moved from X/OPEN UNIX extension to BASE.

23329 Large File Summit extensions are added. In particular, the description of [EINVAL] is clarified  
 23330 and moved from optional to mandatory status.

23331 A note is added to the DESCRIPTION indicating the effects of attempting to lock a section of a  
 23332 file that is associated with a buffered stream.

23333 **Issue 6**

23334 The DESCRIPTION is updated to avoid use of the term “must” for application requirements.

## 23335 NAME

23336 log, logf, logl — natural logarithm function

## 23337 SYNOPSIS

23338 #include &lt;math.h&gt;

23339 double log(double x);

23340 float logf(float x);

23341 long double logl(long double x);

## 23342 DESCRIPTION

23343 CX The functionality described on this reference page is aligned with the ISO C standard. Any  
 23344 conflict between the requirements described here and the ISO C standard is unintentional. This  
 23345 volume of IEEE Std 1003.1-200x defers to the ISO C standard.

23346 These functions shall compute the natural logarithm of their argument  $x$ ,  $\log_e(x)$ .

23347 An application wishing to check for error situations should set *errno* to zero and call  
 23348 *feclearexcept*(FE\_ALL\_EXCEPT) before calling these functions. On return, if *errno* is non-zero or  
 23349 *fetestexcept*(FE\_INVALID | FE\_DIVBYZERO | FE\_OVERFLOW | FE\_UNDERFLOW) is non-  
 23350 zero, an error has occurred.

## 23351 RETURN VALUE

23352 Upon successful completion, these functions shall return the natural logarithm of  $x$ .

23353 If  $x$  is  $\pm 0$ , a pole error shall occur and *log()*, *logf()*, and *logl()* shall return  $-\text{HUGE\_VAL}$ ,  
 23354  $-\text{HUGE\_VALF}$ , and  $-\text{HUGE\_VALL}$ , respectively.

23355 MX For finite values of  $x$  that are less than 0, or if  $x$  is  $-\text{Inf}$ , a domain error shall occur, and either a  
 23356 NaN (if supported), or an implementation-defined value shall be returned.

23357 MX If  $x$  is NaN, a NaN shall be returned.23358 If  $x$  is 1,  $+0$  shall be returned.23359 If  $x$  is  $+\text{Inf}$ ,  $x$  shall be returned.

## 23360 ERRORS

23361 These functions shall fail if:

23362 MX Domain Error The finite value of  $x$  is negative, or  $x$  is  $-\text{Inf}$ .

23363 If the integer expression (math\_errhandling & MATH\_ERRNO) is non-zero, |  
 23364 then *errno* shall be set to [EDOM]. If the integer expression (math\_errhandling |  
 23365 & MATH\_ERREXCEPT) is non-zero, then the invalid floating-point exception |  
 23366 shall be raised. |

23367 Pole Error The value of  $x$  is zero.

23368 If the integer expression (math\_errhandling & MATH\_ERRNO) is non-zero, |  
 23369 then *errno* shall be set to [ERANGE]. If the integer expression |  
 23370 (math\_errhandling & MATH\_ERREXCEPT) is non-zero, then the divide-by- |  
 23371 zero floating-point exception shall be raised. |

23372 **EXAMPLES**

23373 None.

23374 **APPLICATION USAGE**

23375 On error, the expressions (math\_errhandling & MATH\_ERRNO) and (math\_errhandling &  
23376 MATH\_ERREXCEPT) are independent of each other, but at least one of them must be non-zero.

23377 **RATIONALE**

23378 None.

23379 **FUTURE DIRECTIONS**

23380 None.

23381 **SEE ALSO**

23382 *exp()*, *feclearexcept()*, *fetetestexcept()*, *isnan()*, *log10()*, *log1p()*, the Base Definitions volume of |  
23383 IEEE Std 1003.1-200x, Section 4.18, Treatment of Error Conditions for Mathematical Functions, |  
23384 <math.h>

23385 **CHANGE HISTORY**

23386 First released in Issue 1. Derived from Issue 1 of the SVID.

23387 **Issue 5**

23388 The DESCRIPTION is updated to indicate how an application should check for an error. This  
23389 text was previously published in the APPLICATION USAGE section.

23390 **Issue 6**

23391 The DESCRIPTION is updated to avoid use of the term “must” for application requirements.

23392 The *logf()* and *logl()* functions are added for alignment with the ISO/IEC 9899:1999 standard.

23393 The DESCRIPTION, RETURN VALUE, ERRORS, and APPLICATION USAGE sections are  
23394 revised to align with the ISO/IEC 9899:1999 standard.

23395 IEC 60559:1989 standard floating-point extensions over the ISO/IEC 9899:1999 standard are  
23396 marked.

23397 **NAME**

23398 log10, log10f, log10l — base 10 logarithm function

23399 **SYNOPSIS**

23400 #include &lt;math.h&gt;

23401 double log10(double x);

23402 float log10f(float x);

23403 long double log10l(long double x);

23404 **DESCRIPTION**

23405 **CX** The functionality described on this reference page is aligned with the ISO C standard. Any  
 23406 conflict between the requirements described here and the ISO C standard is unintentional. This  
 23407 volume of IEEE Std 1003.1-200x defers to the ISO C standard.

23408 These functions shall compute the base 10 logarithm of their argument  $x$ ,  $\log_{10}(x)$ .

23409 An application wishing to check for error situations should set *errno* to zero and call  
 23410 *feclearexcept*(FE\_ALL\_EXCEPT) before calling these functions. On return, if *errno* is non-zero or  
 23411 *fetestexcept*(FE\_INVALID | FE\_DIVBYZERO | FE\_OVERFLOW | FE\_UNDERFLOW) is non-  
 23412 zero, an error has occurred.

23413 **RETURN VALUE**23414 Upon successful completion, these functions shall return the base 10 logarithm of  $x$ .

23415 If  $x$  is  $\pm 0$ , a pole error shall occur and *log10()*, *log10f()*, and *log10l()* shall return  $-\text{HUGE\_VAL}$ ,  
 23416  $-\text{HUGE\_VALF}$ , and  $-\text{HUGE\_VALL}$ , respectively.

23417 **MX** For finite values of  $x$  that are less than 0, or if  $x$  is  $-\text{Inf}$ , a domain error shall occur, and either a  
 23418 NaN (if supported), or an implementation-defined value shall be returned.

23419 **MX** If  $x$  is NaN, a NaN shall be returned.23420 If  $x$  is 1,  $+0$  shall be returned.23421 If  $x$  is  $+\text{Inf}$ ,  $+\text{Inf}$  shall be returned.23422 **ERRORS**

23423 These functions shall fail if:

23424 **MX** Domain Error The finite value of  $x$  is negative, or  $x$  is  $-\text{Inf}$ .

23425 If the integer expression (math\_errhandling & MATH\_ERRNO) is non-zero, |  
 23426 then *errno* shall be set to [EDOM]. If the integer expression (math\_errhandling |  
 23427 & MATH\_ERREXCEPT) is non-zero, then the invalid floating-point exception |  
 23428 shall be raised. |

23429 Pole Error The value of  $x$  is zero.

23430 If the integer expression (math\_errhandling & MATH\_ERRNO) is non-zero, |  
 23431 then *errno* shall be set to [ERANGE]. If the integer expression |  
 23432 (math\_errhandling & MATH\_ERREXCEPT) is non-zero, then the divide-by- |  
 23433 zero floating-point exception shall be raised. |



23434 **EXAMPLES**

23435 None.

23436 **APPLICATION USAGE**

23437 On error, the expressions (math\_errhandling & MATH\_ERRNO) and (math\_errhandling &  
23438 MATH\_ERREXCEPT) are independent of each other, but at least one of them must be non-zero.

23439 **RATIONALE**

23440 None.

23441 **FUTURE DIRECTIONS**

23442 None.

23443 **SEE ALSO**

23444 *feclearexcept()*, *fetetestexcept()*, *isnan()*, *log()*, *pow()*, the Base Definitions volume of |  
23445 IEEE Std 1003.1-200x, Section 4.18, Treatment of Error Conditions for Mathematical Functions, |  
23446 <math.h>

23447 **CHANGE HISTORY**

23448 First released in Issue 1. Derived from Issue 1 of the SVID.

23449 **Issue 5**

23450 The DESCRIPTION is updated to indicate how an application should check for an error. This  
23451 text was previously published in the APPLICATION USAGE section.

23452 **Issue 6**

23453 The DESCRIPTION is updated to avoid use of the term “must” for application requirements.

23454 The *log10f()* and *log10l()* functions are added for alignment with the ISO/IEC 9899:1999  
23455 standard.

23456 The DESCRIPTION, RETURN VALUE, ERRORS, and APPLICATION USAGE sections are  
23457 revised to align with the ISO/IEC 9899:1999 standard.

23458 IEC 60559:1989 standard floating-point extensions over the ISO/IEC 9899:1999 standard are  
23459 marked.

## 23460 NAME

23461 log1p, log1pf, log1pl — compute a natural logarithm

## 23462 SYNOPSIS

23463 #include &lt;math.h&gt;

23464 double log1p(double x);

23465 float log1pf(float x);

23466 long double log1pl(long double x);

## 23467 DESCRIPTION

23468 CX The functionality described on this reference page is aligned with the ISO C standard. Any  
 23469 conflict between the requirements described here and the ISO C standard is unintentional. This  
 23470 volume of IEEE Std 1003.1-200x defers to the ISO C standard.

23471 These functions shall compute  $\log_e(1.0 + x)$ .

23472 An application wishing to check for error situations should set *errno* to zero and call  
 23473 *feclearexcept*(FE\_ALL\_EXCEPT) before calling these functions. On return, if *errno* is non-zero or  
 23474 *fetestexcept*(FE\_INVALID | FE\_DIVBYZERO | FE\_OVERFLOW | FE\_UNDERFLOW) is non-  
 23475 zero, an error has occurred.

## 23476 RETURN VALUE

23477 Upon successful completion, these functions shall return the natural logarithm of  $1.0 + x$ .

23478 If  $x$  is  $-1$ , a pole error shall occur and *log1p()*, *log1pf()*, and *log1pl()* shall return  $-\text{HUGE\_VAL}$ ,  
 23479  $-\text{HUGE\_VALF}$ , and  $-\text{HUGE\_VALL}$ , respectively.

23480 MX For finite values of  $x$  that are less than  $-1$ , or if  $x$  is  $-\text{Inf}$ , a domain error shall occur, and either a  
 23481 NaN (if supported), or an implementation-defined value shall be returned.

23482 MX If  $x$  is NaN, a NaN shall be returned.23483 If  $x$  is  $\pm 0$ , or  $+\text{Inf}$ ,  $x$  shall be returned.23484 If  $x$  is subnormal, a range error may occur and  $x$  should be returned.

## 23485 ERRORS

23486 These functions shall fail if:

23487 MX Domain Error The finite value of  $x$  is less than  $-1$ , or  $x$  is  $-\text{Inf}$ .

23488 If the integer expression (math\_errhandling & MATH\_ERRNO) is non-zero, |  
 23489 then *errno* shall be set to [EDOM]. If the integer expression (math\_errhandling |  
 23490 & MATH\_ERREXCEPT) is non-zero, then the invalid floating-point exception |  
 23491 shall be raised. |

23492 Pole Error The value of  $x$  is  $-1$ .

23493 If the integer expression (math\_errhandling & MATH\_ERRNO) is non-zero, |  
 23494 then *errno* shall be set to [ERANGE]. If the integer expression |  
 23495 (math\_errhandling & MATH\_ERREXCEPT) is non-zero, then the divide-by- |  
 23496 zero floating-point exception shall be raised. |

23497 These functions may fail if:

23498 MX Range Error The value of  $x$  is subnormal.

23499 If the integer expression (math\_errhandling & MATH\_ERRNO) is non-zero, |  
 23500 then *errno* shall be set to [ERANGE]. If the integer expression |  
 23501 (math\_errhandling & MATH\_ERREXCEPT) is non-zero, then the underflow |  
 23502 floating-point exception shall be raised. |

23503 **EXAMPLES**

23504 None.

23505 **APPLICATION USAGE**

23506 On error, the expressions (math\_errhandling & MATH\_ERRNO) and (math\_errhandling &  
23507 MATH\_ERREXCEPT) are independent of each other, but at least one of them must be non-zero.

23508 **RATIONALE**

23509 None.

23510 **FUTURE DIRECTIONS**

23511 None.

23512 **SEE ALSO**

23513 *feclearexcept()*, *fetestexcept()*, *log()*, the Base Definitions volume of IEEE Std 1003.1-200x, Section |  
23514 4.18, Treatment of Error Conditions for Mathematical Functions, <math.h> |

23515 **CHANGE HISTORY**

23516 First released in Issue 4, Version 2.

23517 **Issue 5**

23518 Moved from X/OPEN UNIX extension to BASE.

23519 **Issue 6**

23520 The DESCRIPTION is updated to avoid use of the term “must” for application requirements.

23521 The *log1p()* function is no longer marked as an extension.

23522 The *log1pf()* and *log1pl()* functions are added for alignment with the ISO/IEC 9899:1999  
23523 standard.

23524 The DESCRIPTION, RETURN VALUE, ERRORS, and APPLICATION USAGE sections are  
23525 revised to align with the ISO/IEC 9899:1999 standard.

23526 IEC 60559:1989 standard floating-point extensions over the ISO/IEC 9899:1999 standard are  
23527 marked.

## 23528 NAME

23529 log2, log2f, log2l — compute base 2 logarithm functions

## 23530 SYNOPSIS

23531 #include &lt;math.h&gt;

23532 double log2(double x);

23533 float log2f(float x);

23534 long double log2l(long double x);

## 23535 DESCRIPTION

23536 CX The functionality described on this reference page is aligned with the ISO C standard. Any  
 23537 conflict between the requirements described here and the ISO C standard is unintentional. This  
 23538 volume of IEEE Std 1003.1-200x defers to the ISO C standard.

23539 These functions shall compute the base 2 logarithm of their argument  $x$ ,  $\log_2(x)$ .

23540 An application wishing to check for error situations should set *errno* to zero and call  
 23541 *feclearexcept*(FE\_ALL\_EXCEPT) before calling these functions. On return, if *errno* is non-zero or  
 23542 *fetestexcept*(FE\_INVALID | FE\_DIVBYZERO | FE\_OVERFLOW | FE\_UNDERFLOW) is non-  
 23543 zero, an error has occurred.

## 23544 RETURN VALUE

23545 Upon successful completion, these functions shall return the base 2 logarithm of  $x$ .

23546 If  $x$  is  $\pm 0$ , a pole error shall occur and *log2()*, *log2f()*, and *log2l()* shall return  $-\text{HUGE\_VAL}$ ,  
 23547  $-\text{HUGE\_VALF}$ , and  $-\text{HUGE\_VALL}$ , respectively.

23548 MX For finite values of  $x$  that are less than 0, or if  $x$  is  $-\text{Inf}$  a domain error shall occur, and either a  
 23549 NaN (if supported), or an implementation-defined value shall be returned.

23550 MX If  $x$  is NaN, a NaN shall be returned.

23551 If  $x$  is 1,  $+0$  shall be returned.

23552 If  $x$  is  $+\text{Inf}$ ,  $x$  shall be returned.

## 23553 ERRORS

23554 These functions shall fail if:

23555 MX Domain Error The finite value of  $x$  is less than zero, or  $x$  is  $-\text{Inf}$ .

23556 If the integer expression (math\_errhandling & MATH\_ERRNO) is non-zero, |  
 23557 then *errno* shall be set to [EDOM]. If the integer expression (math\_errhandling |  
 23558 & MATH\_ERREXCEPT) is non-zero, then the invalid floating-point exception |  
 23559 shall be raised. |

23560 Pole Error The value of  $x$  is zero.

23561 If the integer expression (math\_errhandling & MATH\_ERRNO) is non-zero, |  
 23562 then *errno* shall be set to [ERANGE]. If the integer expression |  
 23563 (math\_errhandling & MATH\_ERREXCEPT) is non-zero, then the divide-by- |  
 23564 zero floating-point exception shall be raised. |

23565 **EXAMPLES**

23566           None.

23567 **APPLICATION USAGE**

23568           On error, the expressions (math\_errhandling & MATH\_ERRNO) and (math\_errhandling &  
23569           MATH\_ERREXCEPT) are independent of each other, but at least one of them must be non-zero.

23570 **RATIONALE**

23571           None.

23572 **FUTURE DIRECTIONS**

23573           None.

23574 **SEE ALSO**

23575           *feclearexcept()*, *fetestexcept()*, *log()*, the Base Definitions volume of IEEE Std 1003.1-200x, Section |  
23576           4.18, Treatment of Error Conditions for Mathematical Functions, <math.h> |

23577 **CHANGE HISTORY**

23578           First released in Issue 6. Derived from the ISO/IEC 9899:1999 standard.

23579 **NAME**

23580 logb, logbf, logbl — radix-independent exponent

23581 **SYNOPSIS**

23582 #include &lt;math.h&gt;

23583 double logb(double x);

23584 float logbf(float x);

23585 long double logbl(long double x);

23586 **DESCRIPTION**

23587 **CX** The functionality described on this reference page is aligned with the ISO C standard. Any  
 23588 conflict between the requirements described here and the ISO C standard is unintentional. This  
 23589 volume of IEEE Std 1003.1-200x defers to the ISO C standard.

23590 These functions shall compute the exponent of  $x$ , which is the integral part of  $\log_r |x|$ , as a  
 23591 signed floating-point value, for non-zero  $x$ , where  $r$  is the radix of the machine's floating-point  
 23592 arithmetic, which is the value of FLT\_RADIX defined in the <float.h> header.

23593 If  $x$  is subnormal it is treated as though it were normalized; thus for finite positive  $x$ :

23594  $1 \leq x * FLT\_RADIX^{-\logb(x)} < FLT\_RADIX$ 

23595 An application wishing to check for error situations should set *errno* to zero and call  
 23596 *feclearexcept*(FE\_ALL\_EXCEPT) before calling these functions. On return, if *errno* is non-zero or  
 23597 *fetestexcept*(FE\_INVALID | FE\_DIVBYZERO | FE\_OVERFLOW | FE\_UNDERFLOW) is non-  
 23598 zero, an error has occurred.

23599 **RETURN VALUE**23600 Upon successful completion, these functions shall return the exponent of  $x$ .

23601 If  $x$  is  $\pm 0$ , a pole error shall occur and *logb*(), *logbf*(), and *logbl*() shall return  $-\text{HUGE\_VAL}$ ,  
 23602  $-\text{HUGE\_VALF}$ , and  $-\text{HUGE\_VALL}$ , respectively.

23603 **MX** If  $x$  is NaN, a NaN shall be returned.23604 If  $x$  is  $\pm\text{Inf}$ ,  $+\text{Inf}$  shall be returned.23605 **ERRORS**

23606 These functions shall fail if:

23607 Pole Error The value of  $x$  is  $\pm 0$ .

23608 If the integer expression (math\_errhandling & MATH\_ERRNO) is non-zero, |  
 23609 then *errno* shall be set to [ERANGE]. If the integer expression |  
 23610 (math\_errhandling & MATH\_ERREXCEPT) is non-zero, then the divide-by- |  
 23611 zero floating-point exception shall be raised. |

23612 **EXAMPLES**

23613 None.

23614 **APPLICATION USAGE**

23615 On error, the expressions (math\_errhandling & MATH\_ERRNO) and (math\_errhandling &  
 23616 MATH\_ERREXCEPT) are independent of each other, but at least one of them must be non-zero.

23617 **RATIONALE**

23618 None.

23619 **FUTURE DIRECTIONS**

23620 None.

23621 **SEE ALSO**

23622 *feclearexcept()*, *fetestexcept()*, *ilogb()*, *scalb()*, the Base Definitions volume of IEEE Std 1003.1-200x, |  
23623 Section 4.18, Treatment of Error Conditions for Mathematical Functions, <float.h>, <math.h> |

23624 **CHANGE HISTORY**

23625 First released in Issue 4, Version 2.

23626 **Issue 5**

23627 Moved from X/OPEN UNIX extension to BASE.

23628 **Issue 6**23629 The *logb()* function is no longer marked as an extension.23630 The *logbf()* and *logbl()* functions are added for alignment with the ISO/IEC 9899:1999 standard.

23631 The DESCRIPTION, RETURN VALUE, ERRORS, and APPLICATION USAGE sections are  
23632 revised to align with the ISO/IEC 9899:1999 standard.

23633 IEC 60559:1989 standard floating-point extensions over the ISO/IEC 9899:1999 standard are  
23634 marked.

23635 **NAME**

23636       logf — natural logarithm function

23637 **SYNOPSIS**

23638       #include <math.h>

23639       float logf(float x);

23640 **DESCRIPTION**

23641       Refer to *log()*.



23642 **NAME**

23643       logl — natural logarithm function

23644 **SYNOPSIS**

23645       #include &lt;math.h&gt;

23646       long double logl(long double x);

23647 **DESCRIPTION**23648       Refer to *log()*.

23649 **NAME**

23650 longjmp — non-local goto

23651 **SYNOPSIS**

23652 #include &lt;setjmp.h&gt;

23653 void longjmp(jmp\_buf env, int val);

23654 **DESCRIPTION**

23655 cx The functionality described on this reference page is aligned with the ISO C standard. Any  
 23656 conflict between the requirements described here and the ISO C standard is unintentional. This  
 23657 volume of IEEE Std 1003.1-200x defers to the ISO C standard.

23658 The *longjmp()* function shall restore the environment saved by the most recent invocation of  
 23659 *setjmp()* in the same thread, with the corresponding **jmp\_buf** argument. If there is no such  
 23660 invocation, or if the function containing the invocation of *setjmp()* has terminated execution in  
 23661 the interim, or if the invocation of *setjmp()* was within the scope of an identifier with variably  
 23662 cx modified type and execution has left that scope in the interim, the behavior is undefined. It is  
 23663 unspecified whether *longjmp()* restores the signal mask, leaves the signal mask unchanged, or  
 23664 restores it to its value at the time *setjmp()* was called.

23665 All accessible objects have values, and all other components of the abstract machine have state  
 23666 (for example, floating-point status flags and open files), as of the time *longjmp()* was called,  
 23667 except that the values of objects of automatic storage duration are unspecified if they meet all  
 23668 the following conditions:

- 23669 • They are local to the function containing the corresponding *setjmp()* invocation.
- 23670 • They do not have volatile-qualified type.
- 23671 • They are changed between the *setjmp()* invocation and *longjmp()* call.

23672 cx As it bypasses the usual function call and return mechanisms, *longjmp()* shall execute correctly  
 23673 in contexts of interrupts, signals, and any of their associated functions. However, if *longjmp()* is  
 23674 invoked from a nested signal handler (that is, from a function invoked as a result of a signal  
 23675 raised during the handling of another signal), the behavior is undefined.

23676 The effect of a call to *longjmp()* where initialization of the **jmp\_buf** structure was not performed  
 23677 in the calling thread is undefined.

23678 **RETURN VALUE**

23679 After *longjmp()* is completed, program execution continues as if the corresponding invocation of  
 23680 *setjmp()* had just returned the value specified by *val*. The *longjmp()* function shall not cause  
 23681 *setjmp()* to return 0; if *val* is 0, *setjmp()* shall return 1.

23682 **ERRORS**

23683 No errors are defined.

23684 **EXAMPLES**

23685 None.

23686 **APPLICATION USAGE**

23687 Applications whose behavior depends on the value of the signal mask should not use *longjmp()*  
 23688 and *setjmp()*, since their effect on the signal mask is unspecified, but should instead use the  
 23689 *siglongjmp()* and *sigsetjmp()* functions (which can save and restore the signal mask under  
 23690 application control).

23691 **RATIONALE**

23692       None.

23693 **FUTURE DIRECTIONS**

23694       None.

23695 **SEE ALSO**

23696       *setjmp()*, *sigaction()*, *siglongjmp()*, *sigsetjmp()*, the Base Definitions volume of  
23697       IEEE Std 1003.1-200x, <**setjmp.h**>

23698 **CHANGE HISTORY**

23699       First released in Issue 1. Derived from Issue 1 of the SVID.

23700 **Issue 5**

23701       The DESCRIPTION is updated for alignment with the POSIX Threads Extension.

23702 **Issue 6**

23703       Extensions beyond the ISO C standard are now marked.

23704       The following new requirements on POSIX implementations derive from alignment with the  
23705       Single UNIX Specification:

- 23706       • The DESCRIPTION now explicitly makes *longjmp()*'s effect on the signal mask unspecified.

23707       The DESCRIPTION is updated for alignment with the ISO/IEC 9899:1999 standard.

23708 **NAME**

23709       lrand48 — generate uniformly distributed pseudo-random non-negative long integers

23710 **SYNOPSIS**

23711 xSI     #include <stdlib.h>

23712       long lrand48(void);

23713

23714 **DESCRIPTION**

23715       Refer to *drand48()*.

23716 **NAME**

23717 lrint, lrintf, lrintl — round to nearest integer value using current rounding direction

23718 **SYNOPSIS**

23719 #include &lt;math.h&gt;

23720 long lrint(double x);

23721 long lrintf(float x);

23722 long lrintl(long double x);

23723 **DESCRIPTION**

23724 **CX** The functionality described on this reference page is aligned with the ISO C standard. Any  
 23725 conflict between the requirements described here and the ISO C standard is unintentional. This  
 23726 volume of IEEE Std 1003.1-200x defers to the ISO C standard.

23727 These functions shall round their argument to the nearest integer value, rounding according to  
 23728 the current rounding direction.

23729 An application wishing to check for error situations should set *errno* to zero and call  
 23730 *feclearexcept*(FE\_ALL\_EXCEPT) before calling these functions. On return, if *errno* is non-zero or  
 23731 *fetestexcept*(FE\_INVALID | FE\_DIVBYZERO | FE\_OVERFLOW | FE\_UNDERFLOW) is non-  
 23732 zero, an error has occurred.

23733 **RETURN VALUE**

23734 Upon successful completion, these functions shall return the rounded integer value.

23735 **MX** If *x* is NaN, a domain error shall occur, and an unspecified value is returned.23736 If *x* is +Inf, a domain error shall occur and an unspecified value is returned.23737 If *x* is -Inf, a domain error shall occur and an unspecified value is returned.

23738 If the correct value is positive and too large to represent as a **long**, a domain error shall occur |  
 23739 and an unspecified value is returned.

23740 If the correct value is negative and too large to represent as a **long**, a domain error shall occur |  
 23741 and an unspecified value is returned.

23742 **ERRORS**

23743 These functions shall fail if:

23744 **MX** Domain Error The *x* argument is NaN or ±Inf, or the correct value is not representable as an  
 23745 integer.

23746 If the integer expression (math\_errhandling & MATH\_ERRNO) is non-zero, |  
 23747 then *errno* shall be set to [EDOM]. If the integer expression (math\_errhandling  
 23748 & MATH\_ERREXCEPT) is non-zero, then the invalid floating-point exception  
 23749 shall be raised. |

23750 **EXAMPLES**

23751 None.

23752 **APPLICATION USAGE**

23753 On error, the expressions (math\_errhandling & MATH\_ERRNO) and (math\_errhandling &  
 23754 MATH\_ERREXCEPT) are independent of each other, but at least one of them must be non-zero.

23755 **RATIONALE**

23756 These functions provide floating-to-integer conversions. They round according to the current  
 23757 rounding direction. If the rounded value is outside the range of the return type, the numeric  
 23758 result is unspecified and the invalid floating-point exception is raised. When they raise no other  
 23759 floating-point exception and the result differs from the argument, they raise the inexact

23760 floating-point exception.

23761 **FUTURE DIRECTIONS**

23762 None.

23763 **SEE ALSO**

23764 *feclearexcept()*, *fetetestexcept()*, *llrint()*, the Base Definitions volume of IEEE Std 1003.1-200x, |

23765 Section 4.18, Treatment of Error Conditions for Mathematical Functions, <math.h> |

23766 **CHANGE HISTORY**

23767 First released in Issue 6. Derived from the ISO/IEC 9899:1999 standard.

23768 **NAME**

23769 lround, lroundf, lroundl — round to nearest integer value

23770 **SYNOPSIS**

```
23771 #include <math.h>
23772 long lround(double x);
23773 long lroundf(float x);
23774 long lroundl(long double x);
```

23775 **DESCRIPTION**

23776 **CX** The functionality described on this reference page is aligned with the ISO C standard. Any  
 23777 conflict between the requirements described here and the ISO C standard is unintentional. This  
 23778 volume of IEEE Std 1003.1-200x defers to the ISO C standard.

23779 These functions shall round their argument to the nearest integer value, rounding halfway cases  
 23780 away from zero, regardless of the current rounding direction.

23781 An application wishing to check for error situations should set *errno* to zero and call  
 23782 *feclearexcept*(FE\_ALL\_EXCEPT) before calling these functions. On return, if *errno* is non-zero or  
 23783 *fetestexcept*(FE\_INVALID | FE\_DIVBYZERO | FE\_OVERFLOW | FE\_UNDERFLOW) is non-  
 23784 zero, an error has occurred.

23785 **RETURN VALUE**

23786 Upon successful completion, these functions shall return the rounded integer value.

23787 **MX** If *x* is NaN, a domain error shall occur, and an unspecified value is returned.23788 If *x* is +Inf, a domain error shall occur and an unspecified value is returned.23789 If *x* is -Inf, a domain error shall occur and an unspecified value is returned.

23790 If the correct value is positive and too large to represent as a **long**, a domain error shall occur |  
 23791 and an unspecified value is returned.

23792 If the correct value is negative and too large to represent as a **long**, a domain error shall occur |  
 23793 and an unspecified value is returned.

23794 **ERRORS**

23795 These functions shall fail if:

23796 **MX** Domain Error The *x* argument is NaN or ±Inf, or the correct value is not representable as an  
 23797 integer.

23798 If the integer expression (math\_errhandling & MATH\_ERRNO) is non-zero, |  
 23799 then *errno* shall be set to [EDOM]. If the integer expression (math\_errhandling  
 23800 & MATH\_ERREXCEPT) is non-zero, then the invalid floating-point exception  
 23801 shall be raised. |

23802 **EXAMPLES**

23803 None.

23804 **APPLICATION USAGE**

23805 On error, the expressions (math\_errhandling & MATH\_ERRNO) and (math\_errhandling &  
 23806 MATH\_ERREXCEPT) are independent of each other, but at least one of them must be non-zero.

23807 **RATIONALE**

23808 These functions provide floating-to-integer conversions. They round according to the current  
 23809 rounding direction. If the rounded value is outside the range of the return type, the numeric  
 23810 result is unspecified and the invalid floating-point exception is raised. When they raise no other  
 23811 floating-point exception and the result differs from the argument, they raise the inexact

23812 floating-point exception.

23813 These functions differ from the *lrint()* functions in the default rounding direction, with the  
23814 *lround()* functions rounding halfway cases away from zero and needing not to raise the inexact  
23815 floating-point exception for non-integer arguments that round to within the range of the return  
23816 type.

23817 **FUTURE DIRECTIONS**

23818 None.

23819 **SEE ALSO**

23820 *feclearexcept()*, *fetetestexcept()*, *llround()*, the Base Definitions volume of IEEE Std 1003.1-200x, |  
23821 Section 4.18, Treatment of Error Conditions for Mathematical Functions, <math.h> |

23822 **CHANGE HISTORY**

23823 First released in Issue 6. Derived from the ISO/IEC 9899:1999 standard.



23824 **NAME**

23825 lsearch, lfind — linear search and update

23826 **SYNOPSIS**

```
23827 xSI #include <search.h>
23828 void *lsearch(const void *key, void *base, size_t *nel, size_t width,
23829             int (*compar)(const void *, const void *));
23830 void *lfind(const void *key, const void *base, size_t *nel,
23831           size_t width, int (*compar)(const void *, const void *));
23832
```

23833 **DESCRIPTION**

23834 The *lsearch()* function shall linearly search the table and return a pointer into the table for the  
 23835 matching entry. If the entry does not occur, it shall be added at the end of the table. The *key*  
 23836 argument points to the entry to be sought in the table. The *base* argument points to the first  
 23837 element in the table. The *width* argument is the size of an element in bytes. The *nel* argument  
 23838 points to an integer containing the current number of elements in the table. The integer to which  
 23839 *nel* points shall be incremented if the entry is added to the table. The *compar* argument points to  
 23840 a comparison function which the application shall supply (for example, *strcmp()*). It is called  
 23841 with two arguments that point to the elements being compared. The application shall ensure  
 23842 that the function returns 0 if the elements are equal, and non-zero otherwise.

23843 The *lfind()* function shall be equivalent to *lsearch()*, except that if the entry is not found, it is not  
 23844 added to the table. Instead, a null pointer is returned.

23845 **RETURN VALUE**

23846 If the searched for entry is found, both *lsearch()* and *lfind()* shall return a pointer to it. Otherwise,  
 23847 *lfind()* shall return a null pointer and *lsearch()* shall return a pointer to the newly added element.

23848 Both functions shall return a null pointer in case of error.

23849 **ERRORS**

23850 No errors are defined.

23851 **EXAMPLES**23852 **Storing Strings in a Table**

23853 This fragment reads in less than or equal to TABSIZE strings of length less than or equal to  
 23854 ELSIZE and stores them in a table, eliminating duplicates.

```
23855 #include <stdio.h>
23856 #include <string.h>
23857 #include <search.h>
23858 #define TABSIZE 50
23859 #define ELSIZE 120
23860 ...
23861     char line[ELSIZE], tab[TABSIZE][ELSIZE];
23862     size_t nel = 0;
23863     ...
23864     while (fgets(line, ELSIZE, stdin) != NULL && nel < TABSIZE)
23865         (void) lsearch(line, tab, &nel,
23866                       ELSIZE, (int (*)(const void *, const void *)) strcmp);
23867     ...
```

**23868 Finding a Matching Entry**

23869 The following example finds any line that reads "This is a test."

```
23870 #include <search.h>
23871 #include <string.h>
23872 ...
23873 char line[ELSIZE], tab[TABSIZE][ELSIZE];
23874 size_t nel = 0;
23875 char *findline;
23876 void *entry;

23877 findline = "This is a test.\n";

23878 entry = lfind(findline, tab, &nel, ELSIZE, (
23879     int (*)(const void *, const void *)) strcmp);
```

**23880 APPLICATION USAGE**

23881 The comparison function need not compare every byte, so arbitrary data may be contained in  
23882 the elements in addition to the values being compared.

23883 Undefined results can occur if there is not enough room in the table to add a new item.

**23884 RATIONALE**

23885 None.

**23886 FUTURE DIRECTIONS**

23887 None.

**23888 SEE ALSO**

23889 *hcreate()*, *tsearch()*, the Base Definitions volume of IEEE Std 1003.1-200x, <**search.h**>

**23890 CHANGE HISTORY**

23891 First released in Issue 1. Derived from Issue 1 of the SVID.

**23892 Issue 6**

23893 The DESCRIPTION is updated to avoid use of the term “must” for application requirements.

23894 **NAME**

23895 lseek — move the read/write file offset

23896 **SYNOPSIS**

23897 #include &lt;unistd.h&gt;

23898 off\_t lseek(int *fil-des*, off\_t *offset*, int *whence*);23899 **DESCRIPTION**23900 The *lseek()* function shall set the file offset for the open file description associated with the file descriptor *fil-des*, as follows:

- 23902 • If *whence* is SEEK\_SET, the file offset shall be set to *offset* bytes. |
- 23903 • If *whence* is SEEK\_CUR, the file offset shall be set to its current location plus *offset*. |
- 23904 • If *whence* is SEEK\_END, the file offset shall be set to the size of the file plus *offset*. |

23905 The symbolic constants SEEK\_SET, SEEK\_CUR, and SEEK\_END are defined in &lt;unistd.h&gt;.

23906 The behavior of *lseek()* on devices which are incapable of seeking is implementation-defined.  
23907 The value of the file offset associated with such a device is undefined.23908 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.23911 The *lseek()* function shall not, by itself, extend the size of a file.23912 SHM If *fil-des* refers to a shared memory object, the result of the *lseek()* function is unspecified.23913 TYM If *fil-des* refers to a typed memory object, the result of the *lseek()* function is unspecified.23914 **RETURN VALUE**23915 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.23918 **ERRORS**23919 The *lseek()* function shall fail if:

- 23920 [EBADF] The *fil-des* argument is not an open file descriptor.
- 23921 [EINVAL] The *whence* argument is not a proper value, or the resulting file offset would be negative for a regular file, block special file, or directory.
- 23922
- 23923 [EOVERFLOW] The resulting file offset would be a value which cannot be represented correctly in an object of type **off\_t**.
- 23924
- 23925 [ESPIPE] The *fil-des* argument is associated with a pipe, FIFO, or socket.

23926 **EXAMPLES**

23927 None.

23928 **APPLICATION USAGE**

23929 None.

23930 **RATIONALE**23931 The ISO C standard includes the functions *fgetpos()* and *fsetpos()*, which work on very large files by use of a special positioning type.23933 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 IEEE Std 1003.1-200x that may directly |

23934

- 23935 extend the size of the file is *write()*, *truncate()*, and *ftruncate()*, several functions originally |  
23936 derived from the ISO C standard, such as *fwrite()*, *fprintf()*, and so on, may do so (by causing |  
23937 calls on *write()*).
- 23938 An invalid file offset that would cause [EINVAL] to be returned may be both implementation-  
23939 defined and device-dependent (for example, memory may have few invalid values). A negative  
23940 file offset may be valid for some devices in some implementations.
- 23941 The POSIX.1-1990 standard did not specifically prohibit *lseek()* from returning a negative offset.  
23942 Therefore, an application was required to clear *errno* prior to the call and check *errno* upon return  
23943 to determine whether a return value of (*off\_t*)-1 is a negative offset or an indication of an error  
23944 condition. The standard developers did not wish to require this action on the part of a |  
23945 conforming application, and chose to require that *errno* be set to [EINVAL] when the resulting |  
23946 file offset would be negative for a regular file, block special file, or directory.
- 23947 **FUTURE DIRECTIONS**
- 23948 None.
- 23949 **SEE ALSO**
- 23950 *open()*, the Base Definitions volume of IEEE Std 1003.1-200x, <**sys/types.h**>, <**unistd.h**>
- 23951 **CHANGE HISTORY**
- 23952 First released in Issue 1. Derived from Issue 1 of the SVID.
- 23953 **Issue 5**
- 23954 The DESCRIPTION is updated for alignment with the POSIX Realtime Extension.
- 23955 Large File Summit extensions are added.
- 23956 **Issue 6**
- 23957 In the SYNOPSIS, the optional include of the <**sys/types.h**> header is removed.
- 23958 The following new requirements on POSIX implementations derive from alignment with the  
23959 Single UNIX Specification:
- 23960 • The requirement to include <**sys/types.h**> has been removed. Although <**sys/types.h**> was  
23961 required for conforming implementations of previous POSIX specifications, it was not  
23962 required for UNIX applications.
  - 23963 • The [EOVERFLOW] error condition is added. This change is to support large files.
- 23964 An additional [ESPIPE] error condition is added for sockets.
- 23965 The DESCRIPTION is updated for alignment with IEEE Std 1003.1j-2000 by specifying that  
23966 *lseek()* results are unspecified for typed memory objects.

23967 **NAME**

23968        lstat — get symbolic link status

23969 **SYNOPSIS**

23970        #include &lt;sys/stat.h&gt;

23971        int lstat(const char \*restrict *path*, struct stat \*restrict *buf*);23972 **DESCRIPTION**

23973        The *lstat()* function shall be equivalent to *stat()*, except when *path* refers to a symbolic link. In |  
 23974        that case *lstat()* shall return information about the link, while *stat()* shall return information |  
 23975        about the file the link references.

23976        For symbolic links, the *st\_mode* member shall contain meaningful information when used with |  
 23977        the file type macros, and the *st\_size* member shall contain the length of the pathname contained |  
 23978        in the symbolic link. File mode bits and the contents of the remaining members of the **stat** |  
 23979        structure are unspecified. The value returned in the *st\_size* member is the length of the contents |  
 23980        of the symbolic link, and does not count any trailing null.

23981 **RETURN VALUE**

23982        Upon successful completion, *lstat()* shall return 0. Otherwise, it shall return -1 and set *errno* to |  
 23983        indicate the error.

23984 **ERRORS**23985        The *lstat()* function shall fail if:

23986        [EACCES]        A component of the path prefix denies search permission.

23987        [EIO]            An error occurred while reading from the file system.

23988        [ELOOP]         A loop exists in symbolic links encountered during resolution of the *path* |  
 23989        argument.

23990        [ENAMETOOLONG]

23991                        The length of a pathname exceeds {PATH\_MAX} or a pathname component is |  
 23992                        longer than {NAME\_MAX}. |

23993        [ENOTDIR]        A component of the path prefix is not a directory.

23994        [ENOENT]         A component of *path* does not name an existing file or *path* is an empty string.23995        [EOVERFLOW]     The file size in bytes or the number of blocks allocated to the file or the file |  
 23996                        serial number cannot be represented correctly in the structure pointed to by |  
 23997                        *buf*.23998        The *lstat()* function may fail if:23999        [ELOOP]         More than {SYMLOOP\_MAX} symbolic links were encountered during |  
 24000                        resolution of the *path* argument.

24001        [ENAMETOOLONG]

24002                        As a result of encountering a symbolic link in resolution of the *path* argument, |  
 24003                        the length of the substituted pathname string exceeded {PATH\_MAX}. |24004        [EOVERFLOW]     One of the members is too large to store into the structure pointed to by the |  
 24005                        *buf* argument.

24006 **EXAMPLES**24007 **Obtaining Symbolic Link Status Information**

24008 The following example shows how to obtain status information for a symbolic link named  
24009 **/modules/pass1**. The structure variable *buffer* is defined for the **stat** structure. If the *path*  
24010 argument specified the filename for the file pointed to by the symbolic link (**/home/cnd/mod1**),  
24011 the results of calling the function would be the same as those returned by a call to the *stat()*  
24012 function.

```
24013 #include <sys/stat.h>
24014 struct stat buffer;
24015 int status;
24016 ...
24017 status = lstat("/modules/pass1", &buffer);
```

24018 **APPLICATION USAGE**

24019 None.

24020 **RATIONALE**

24021 The *lstat()* function is not required to update the time-related fields if the named file is not a  
24022 symbolic link. While the *st\_uid*, *st\_gid*, *st\_atime*, *st\_mtime*, and *st\_ctime* members of the **stat**  
24023 structure may apply to a symbolic link, they are not required to do so. No functions in  
24024 IEEE Std 1003.1-200x are required to maintain any of these time fields.

24025 **FUTURE DIRECTIONS**

24026 None.

24027 **SEE ALSO**

24028 *lstat()*, *readlink()*, *stat()*, *symlink()*, the Base Definitions volume of IEEE Std 1003.1-200x,  
24029 **<sys/stat.h>**

24030 **CHANGE HISTORY**

24031 First released in Issue 4, Version 2.

24032 **Issue 5**

24033 Moved from X/OPEN UNIX extension to BASE.

24034 Large File Summit extensions are added.

24035 **Issue 6**

24036 The following changes were made to align with the IEEE P1003.1a draft standard:

- 24037 • This function is now mandatory.
- 24038 • The [ELOOP] optional error condition is added.

24039 The **restrict** keyword is added to the *lstat()* prototype for alignment with the ISO/IEC 9899:1999  
24040 standard.

24041 **NAME**

24042       makecontext, swapcontext — manipulate user contexts

24043 **SYNOPSIS**

```
24044 xSI      #include <ucontext.h>
24045          void makecontext(ucontext_t *ucp, void (*func)(void),
24046                          int argc, ...);
24047          int swapcontext(ucontext_t *restrict oucp,
24048                          const ucontext_t *restrict ucp);
24049
```

24050 **DESCRIPTION**

24051       The *makecontext()* function shall modify the context specified by *ucp*, which has been initialized  
 24052       using *getcontext()*. When this context is resumed using *swapcontext()* or *setcontext()*, program  
 24053       execution shall continue by calling *func*, passing it the arguments that follow *argc* in the  
 24054       *makecontext()* call.

24055       Before a call is made to *makecontext()*, the application shall ensure that the context being  
 24056       modified has a stack allocated for it. The application shall ensure that the value of *argc* matches  
 24057       the number of integer arguments passed to *func*; otherwise, the behavior is undefined.

24058       The *uc\_link* member is used to determine the context that shall be resumed when the context  
 24059       being modified by *makecontext()* returns. The application shall ensure that the *uc\_link* member is  
 24060       initialized prior to the call to *makecontext()*.

24061       The *swapcontext()* function shall save the current context in the context structure pointed to by  
 24062       *oucp* and shall set the context to the context structure pointed to by *ucp*.

24063 **RETURN VALUE**

24064       Upon successful completion, *swapcontext()* shall return 0. Otherwise,  $-1$  shall be returned and  
 24065       *errno* set to indicate the error.

24066 **ERRORS**

24067       The *swapcontext()* function shall fail if:

24068       [ENOMEM]       The *ucp* argument does not have enough stack left to complete the operation.

24069 **EXAMPLES**

24070       The following example illustrates the use of *makecontext()*:

```
24071      #include <stdio.h>
24072      #include <ucontext.h>
24073
24074      static ucontext_t ctx[3];
24075
24076      static void
24077      f1 (void)
24078      {
24079          puts("start f1");
24080          swapcontext(&ctx[1], &ctx[2]);
24081          puts("finish f1");
24082      }
24083
24084      static void
24085      f2 (void)
24086      {
24087          puts("start f2");
24088          swapcontext(&ctx[2], &ctx[1]);
24089      }
```

```

24086         puts("finish f2");
24087     }
24088     int
24089     main (void)
24090     {
24091         char st1[8192];
24092         char st2[8192];
24093
24094         getcontext(&ctx[1]);
24095         ctx[1].uc_stack.ss_sp = st1;
24096         ctx[1].uc_stack.ss_size = sizeof st1;
24097         ctx[1].uc_link = &ctx[0];
24098         makecontext(&ctx[1], f1, 0);
24099
24100         getcontext(&ctx[2]);
24101         ctx[2].uc_stack.ss_sp = st2;
24102         ctx[2].uc_stack.ss_size = sizeof st2;
24103         ctx[2].uc_link = &ctx[1];
24104         makecontext(&ctx[2], f2, 0);
24105
24106         swapcontext(&ctx[0], &ctx[2]);
24107         return 0;
24108     }

```

#### 24106 APPLICATION USAGE

24107 None.

#### 24108 RATIONALE

24109 None.

#### 24110 FUTURE DIRECTIONS

24111 None.

#### 24112 SEE ALSO

24113 *exit()*, *getcontext()*, *sigaction()*, *sigprocmask()*, the Base Definitions volume of  
 24114 IEEE Std 1003.1-200x, <**ucontext.h**>

#### 24115 CHANGE HISTORY

24116 First released in Issue 4, Version 2.

#### 24117 Issue 5

24118 Moved from X/OPEN UNIX extension to BASE.

24119 In the ERRORS section, the description of [ENOMEM] is changed to apply to *swapcontext()* only.

#### 24120 Issue 6

24121 The DESCRIPTION is updated to avoid use of the term “must” for application requirements.

24122 The **restrict** keyword is added to the *swapcontext()* prototype for alignment with the  
 24123 ISO/IEC 9899:1999 standard.



24124 **NAME**

24125 malloc — a memory allocator

24126 **SYNOPSIS**

24127 #include &lt;stdlib.h&gt;

24128 void \*malloc(size\_t size);

24129 **DESCRIPTION**

24130 CX The functionality described on this reference page is aligned with the ISO C standard. Any  
 24131 conflict between the requirements described here and the ISO C standard is unintentional. This  
 24132 volume of IEEE Std 1003.1-200x defers to the ISO C standard.

24133 The *malloc()* function shall allocate unused space for an object whose size in bytes is specified by  
 24134 *size* and whose value is unspecified.

24135 The order and contiguity of storage allocated by successive calls to *malloc()* is unspecified. The  
 24136 pointer returned if the allocation succeeds shall be suitably aligned so that it may be assigned to  
 24137 a pointer to any type of object and then used to access such an object in the space allocated (until  
 24138 the space is explicitly freed or reallocated). Each such allocation shall yield a pointer to an object  
 24139 disjoint from any other object. The pointer returned points to the start (lowest byte address) of  
 24140 the allocated space. If the space cannot be allocated, a null pointer shall be returned. If the size of  
 24141 the space requested is 0, the behavior is implementation-defined: the value returned shall be  
 24142 either a null pointer or a unique pointer.

24143 **RETURN VALUE**

24144 Upon successful completion with *size* not equal to 0, *malloc()* shall return a pointer to the  
 24145 allocated space. If *size* is 0, either a null pointer or a unique pointer that can be successfully  
 24146 CX passed to *free()* shall be returned. Otherwise, it shall return a null pointer and set *errno* to  
 24147 indicate the error.

24148 **ERRORS**24149 The *malloc()* function shall fail if:

24150 CX [ENOMEM] Insufficient storage space is available.

24151 **EXAMPLES**

24152 None.

24153 **APPLICATION USAGE**

24154 None.

24155 **RATIONALE**

24156 None.

24157 **FUTURE DIRECTIONS**

24158 None.

24159 **SEE ALSO**24160 *calloc()*, *free()*, *realloc()*, the Base Definitions volume of IEEE Std 1003.1-200x, <stdlib.h>24161 **CHANGE HISTORY**

24162 First released in Issue 1. Derived from Issue 1 of the SVID.

24163 **Issue 6**

24164 Extensions beyond the ISO C standard are now marked.

24165 The following new requirements on POSIX implementations derive from alignment with the  
 24166 Single UNIX Specification:

24167

- In the RETURN VALUE section, the requirement to set *errno* to indicate an error is added.

24168

- The [ENOMEM] error condition is added.

24169 **NAME**

24170           mblen — get number of bytes in a character

24171 **SYNOPSIS**

24172           #include &lt;stdlib.h&gt;

24173           int mblen(const char \*s, size\_t n);

24174 **DESCRIPTION**

24175 **CX**       The functionality described on this reference page is aligned with the ISO C standard. Any  
 24176 conflict between the requirements described here and the ISO C standard is unintentional. This  
 24177 volume of IEEE Std 1003.1-200x defers to the ISO C standard.

24178       If *s* is not a null pointer, *mblen()* shall determine the number of bytes constituting the character  
 24179 pointed to by *s*. Except that the shift state of *mbtowc()* is not affected, it shall be equivalent to:

24180           mbtowc((wchar\_t \*)0, s, n);

24181       The implementation shall behave as if no function defined in this volume of  
 24182 IEEE Std 1003.1-200x calls *mblen()*.

24183       The behavior of this function is affected by the *LC\_CTYPE* category of the current locale. For a  
 24184 state-dependent encoding, this function shall be placed into its initial state by a call for which its  
 24185 character pointer argument, *s*, is a null pointer. Subsequent calls with *s* as other than a null  
 24186 pointer shall cause the internal state of the function to be altered as necessary. A call with *s* as a  
 24187 null pointer shall cause this function to return a non-zero value if encodings have state  
 24188 dependency, and 0 otherwise. If the implementation employs special bytes to change the shift  
 24189 state, these bytes shall not produce separate wide-character codes, but shall be grouped with an  
 24190 adjacent character. Changing the *LC\_CTYPE* category causes the shift state of this function to be  
 24191 unspecified.

24192 **RETURN VALUE**

24193       If *s* is a null pointer, *mblen()* shall return a non-zero or 0 value, if character encodings,  
 24194 respectively, do or do not have state-dependent encodings. If *s* is not a null pointer, *mblen()* shall  
 24195 either return 0 (if *s* points to the null byte), or return the number of bytes that constitute the  
 24196 character (if the next *n* or fewer bytes form a valid character), or return -1 (if they do not form a  
 24197 **CX** valid character) and may set *errno* to indicate the error. In no case shall the value returned be  
 24198 greater than *n* or the value of the {MB\_CUR\_MAX} macro.

24199 **ERRORS**24200       The *mblen()* function may fail if:24201 **XSI**       [EILSEQ]       Invalid character sequence is detected.24202 **EXAMPLES**

24203       None.

24204 **APPLICATION USAGE**

24205       None.

24206 **RATIONALE**

24207       None.

24208 **FUTURE DIRECTIONS**

24209       None.

24210 **SEE ALSO**

24211 *mbtowc()*, *mbstowcs()*, *wctomb()*, *wcstombs()*, the Base Definitions volume of  
24212 IEEE Std 1003.1-200x, <stdlib.h>

24213 **CHANGE HISTORY**

24214 First released in Issue 4. Aligned with the ISO C standard.

24215 **NAME**24216 `mbrlen` — get number of bytes in a character (restartable)24217 **SYNOPSIS**24218 `#include <wchar.h>`24219 `size_t mbrlen(const char *restrict s, size_t n,`  
24220 `mbstate_t *restrict ps);`24221 **DESCRIPTION**24222 **CX** The functionality described on this reference page is aligned with the ISO C standard. Any  
24223 conflict between the requirements described here and the ISO C standard is unintentional. This  
24224 volume of IEEE Std 1003.1-200x defers to the ISO C standard.24225 If *s* is not a null pointer, `mbrlen()` shall determine the number of bytes constituting the character  
24226 pointed to by *s*. It shall be equivalent to:24227 `mbstate_t internal;`24228 `mbrtowc(NULL, s, n, ps != NULL ? ps : &internal);`24229 If *ps* is a null pointer, the `mbrlen()` function shall use its own internal **mbstate\_t** object, which is  
24230 initialized at program start-up to the initial conversion state. Otherwise, the **mbstate\_t** object  
24231 pointed to by *ps* shall be used to completely describe the current conversion state of the  
24232 associated character sequence. The implementation shall behave as if no function defined in this  
24233 volume of IEEE Std 1003.1-200x calls `mbrlen()`.24234 The behavior of this function is affected by the `LC_CTYPE` category of the current locale.24235 **RETURN VALUE**24236 The `mbrlen()` function shall return the first of the following that applies:24237 **0** If the next *n* or fewer bytes complete the character that corresponds to the null  
24238 wide character.24239 **positive** If the next *n* or fewer bytes complete a valid character; the value returned shall  
24240 be the number of bytes that complete the character.24241 **(size\_t)-2** If the next *n* bytes contribute to an incomplete but potentially valid character,  
24242 and all *n* bytes have been processed. When *n* has at least the value of the  
24243 `{MB_CUR_MAX}` macro, this case can only occur if *s* points at a sequence of  
24244 redundant shift sequences (for implementations with state-dependent  
24245 encodings).24246 **(size\_t)-1** If an encoding error occurs, in which case the next *n* or fewer bytes do not  
24247 contribute to a complete and valid character. In this case, `[EILSEQ]` shall be  
24248 stored in `errno` and the conversion state is undefined.24249 **ERRORS**24250 The `mbrlen()` function may fail if:24251 `[EINVAL]` *ps* points to an object that contains an invalid conversion state.24252 `[EILSEQ]` Invalid character sequence is detected.

24253 **EXAMPLES**

24254 None.

24255 **APPLICATION USAGE**

24256 None.

24257 **RATIONALE**

24258 None.

24259 **FUTURE DIRECTIONS**

24260 None.

24261 **SEE ALSO**24262 *mbsinit()*, *mbrtowc()*, the Base Definitions volume of IEEE Std 1003.1-200x, <**wchar.h**>24263 **CHANGE HISTORY**

24264 First released in Issue 5. Included for alignment with ISO/IEC 9899:1990/Amendment 1:1995

24265 (E).

24266 **Issue 6**24267 The *mbrlen()* prototype is updated for alignment with the ISO/IEC 9899:1999 standard.

24268 **NAME**

24269 mbrtowc — convert a character to a wide-character code (restartable)

24270 **SYNOPSIS**

24271 #include &lt;wchar.h&gt;

24272 size\_t mbrtowc(wchar\_t \*restrict *pwc*, const char \*restrict *s*,  
24273 size\_t *n*, mbstate\_t \*restrict *ps*);24274 **DESCRIPTION**24275 **CX** The functionality described on this reference page is aligned with the ISO C standard. Any  
24276 conflict between the requirements described here and the ISO C standard is unintentional. This  
24277 volume of IEEE Std 1003.1-200x defers to the ISO C standard.24278 If *s* is a null pointer, the *mbrtowc()* function shall be equivalent to the call:24279 mbrtowc(NULL, "", 1, *ps*)24280 In this case, the values of the arguments *pwc* and *n* are ignored.24281 If *s* is not a null pointer, the *mbrtowc()* function shall inspect at most *n* bytes beginning at the  
24282 byte pointed to by *s* to determine the number of bytes needed to complete the next character  
24283 (including any shift sequences). If the function determines that the next character is completed, it  
24284 shall determine the value of the corresponding wide character and then, if *pwc* is not a null  
24285 pointer, shall store that value in the object pointed to by *pwc*. If the corresponding wide  
24286 character is the null wide character, the resulting state described shall be the initial conversion  
24287 state.24288 If *ps* is a null pointer, the *mbrtowc()* function shall use its own internal **mbstate\_t** object, which  
24289 shall be initialized at program start-up to the initial conversion state. Otherwise, the **mbstate\_t**  
24290 object pointed to by *ps* shall be used to completely describe the current conversion state of the  
24291 associated character sequence. The implementation shall behave as if no function defined in this  
24292 volume of IEEE Std 1003.1-200x calls *mbrtowc()*.24293 The behavior of this function is affected by the *LC\_CTYPE* category of the current locale.24294 **RETURN VALUE**24295 The *mbrtowc()* function shall return the first of the following that applies:24296 **0** If the next *n* or fewer bytes complete the character that corresponds to the null  
24297 wide character (which is the value stored).24298 between 1 and *n* inclusive24299 If the next *n* or fewer bytes complete a valid character (which is the value  
24300 stored); the value returned shall be the number of bytes that complete the  
24301 character.24302 **(size\_t)–2** If the next *n* bytes contribute to an incomplete but potentially valid character,  
24303 and all *n* bytes have been processed (no value is stored). When *n* has at least  
24304 the value of the {*MB\_CUR\_MAX*} macro, this case can only occur if *s* points at  
24305 a sequence of redundant shift sequences (for implementations with state-  
24306 dependent encodings).24307 **(size\_t)–1** If an encoding error occurs, in which case the next *n* or fewer bytes do not  
24308 contribute to a complete and valid character (no value is stored). In this case,  
24309 [EILSEQ] shall be stored in *errno* and the conversion state is undefined.

24310 **ERRORS**

24311 The *mbrtowc()* function may fail if:

24312 CX [EINVAL] *ps* points to an object that contains an invalid conversion state.

24313 [EILSEQ] Invalid character sequence is detected.

24314 **EXAMPLES**

24315 None.

24316 **APPLICATION USAGE**

24317 None.

24318 **RATIONALE**

24319 None.

24320 **FUTURE DIRECTIONS**

24321 None.

24322 **SEE ALSO**

24323 *mbstinit()*, the Base Definitions volume of IEEE Std 1003.1-200x, <*wchar.h*>

24324 **CHANGE HISTORY**

24325 First released in Issue 5. Included for alignment with ISO/IEC 9899:1990/Amendment 1:1995  
24326 (E).

24327 **Issue 6**

24328 The *mbrtowc()* prototype is updated for alignment with the ISO/IEC 9899:1999 standard. |

24329 The following new requirements on POSIX implementations derive from alignment with the |  
24330 Single UNIX Specification: |

24331 • The [EINVAL] error condition is added. |

24332 ISO/IEC 9899:1999 standard, Technical Corrigendum No. 1 is incorporated. |



24333 **NAME**

24334 `mbsinit` — determine conversion object status

24335 **SYNOPSIS**

24336 `#include <wchar.h>`

24337 `int mbsinit(const mbstate_t *ps);`

24338 **DESCRIPTION**

24339 **CX** The functionality described on this reference page is aligned with the ISO C standard. Any  
24340 conflict between the requirements described here and the ISO C standard is unintentional. This  
24341 volume of IEEE Std 1003.1-200x defers to the ISO C standard.

24342 If *ps* is not a null pointer, the *mbsinit()* function shall determine whether the object pointed to by  
24343 *ps* describes an initial conversion state.

24344 **RETURN VALUE**

24345 The *mbsinit()* function shall return non-zero if *ps* is a null pointer, or if the pointed-to object  
24346 describes an initial conversion state; otherwise, it shall return zero.

24347 If an **mbstate\_t** object is altered by any of the functions described as “restartable”, and is then  
24348 used with a different character sequence, or in the other conversion direction, or with a different  
24349 *LC\_CTYPE* category setting than on earlier function calls, the behavior is undefined.

24350 **ERRORS**

24351 No errors are defined.

24352 **EXAMPLES**

24353 None.

24354 **APPLICATION USAGE**

24355 The **mbstate\_t** object is used to describe the current conversion state from a particular character  
24356 sequence to a wide-character sequence (or *vice versa*) under the rules of a particular setting of the  
24357 *LC\_CTYPE* category of the current locale.

24358 The initial conversion state corresponds, for a conversion in either direction, to the beginning of  
24359 a new character sequence in the initial shift state. A zero valued **mbstate\_t** object is at least one  
24360 way to describe an initial conversion state. A zero valued **mbstate\_t** object can be used to initiate  
24361 conversion involving any character sequence, in any *LC\_CTYPE* category setting.

24362 **RATIONALE**

24363 None.

24364 **FUTURE DIRECTIONS**

24365 None.

24366 **SEE ALSO**

24367 *mbrlen()*, *mbrtowc()*, *wcrtomb()*, *mbsrtowcs()*, *wcsrtombs()*, the Base Definitions volume of  
24368 IEEE Std 1003.1-200x, `<wchar.h>`

24369 **CHANGE HISTORY**

24370 First released in Issue 5. Included for alignment with ISO/IEC 9899:1990/Amendment 1:1995  
24371 (E).

## 24372 NAME

24373 mbsrtowcs — convert a character string to a wide-character string (restartable)

## 24374 SYNOPSIS

24375 #include <wchar.h>

24376 size\_t mbsrtowcs(wchar\_t \*restrict dst, const char \*\*restrict src,  
24377 size\_t len, mbstate\_t \*restrict ps);

## 24378 DESCRIPTION

24379 CX The functionality described on this reference page is aligned with the ISO C standard. Any  
24380 conflict between the requirements described here and the ISO C standard is unintentional. This  
24381 volume of IEEE Std 1003.1-200x defers to the ISO C standard.

24382 The *mbsrtowcs()* function shall convert a sequence of characters, beginning in the conversion  
24383 state described by the object pointed to by *ps*, from the array indirectly pointed to by *src* into a  
24384 sequence of corresponding wide characters. If *dst* is not a null pointer, the converted characters  
24385 shall be stored into the array pointed to by *dst*. Conversion continues up to and including a  
24386 terminating null character, which shall also be stored. Conversion shall stop early in either of the  
24387 following cases:

- 24388 • A sequence of bytes is encountered that does not form a valid character.
- 24389 • *len* codes have been stored into the array pointed to by *dst* (and *dst* is not a null pointer).

24390 Each conversion shall take place as if by a call to the *mbrtowc()* function.

24391 If *dst* is not a null pointer, the pointer object pointed to by *src* shall be assigned either a null  
24392 pointer (if conversion stopped due to reaching a terminating null character) or the address just  
24393 past the last character converted (if any). If conversion stopped due to reaching a terminating  
24394 null character, and if *dst* is not a null pointer, the resulting state described shall be the initial  
24395 conversion state.

24396 If *ps* is a null pointer, the *mbsrtowcs()* function shall use its own internal **mbstate\_t** object, which  
24397 is initialized at program start-up to the initial conversion state. Otherwise, the **mbstate\_t** object  
24398 pointed to by *ps* shall be used to completely describe the current conversion state of the  
24399 associated character sequence. The implementation behaves as if no function defined in this  
24400 volume of IEEE Std 1003.1-200x calls *mbsrtowcs()*.

24401 The behavior of this function shall be affected by the *LC\_CTYPE* category of the current locale.

## 24402 RETURN VALUE

24403 If the input conversion encounters a sequence of bytes that do not form a valid character, an  
24404 encoding error occurs. In this case, the *mbsrtowcs()* function stores the value of the macro  
24405 [EILSEQ] in *errno* and shall return (**size\_t**)-1; the conversion state is undefined. Otherwise, it  
24406 shall return the number of characters successfully converted, not including the terminating null  
24407 (if any).

## 24408 ERRORS

24409 The *mbsrtowcs()* function may fail if:

- 24410 CX [EINVAL] *ps* points to an object that contains an invalid conversion state.
- 24411 [EILSEQ] Invalid character sequence is detected.

24412 **EXAMPLES**

24413 None.

24414 **APPLICATION USAGE**

24415 None.

24416 **RATIONALE**

24417 None.

24418 **FUTURE DIRECTIONS**

24419 None.

24420 **SEE ALSO**24421 *mbstowcs()*, *mbstowc()*, the Base Definitions volume of IEEE Std 1003.1-200x, <wchar.h>24422 **CHANGE HISTORY**24423 First released in Issue 5. Included for alignment with ISO/IEC 9899:1990/Amendment 1:1995  
24424 (E).24425 **Issue 6**24426 The *mbsrtowcs()* prototype is updated for alignment with the ISO/IEC 9899:1999 standard. |

24427 The [EINVAL] error condition is marked CX. |

24428 **NAME**

24429 mbstowcs — convert a character string to a wide-character string

24430 **SYNOPSIS**

24431 #include &lt;stdlib.h&gt;

24432 size\_t mbstowcs(wchar\_t \*restrict pwcs, const char \*restrict s,  
24433 size\_t n);24434 **DESCRIPTION**24435 **CX** The functionality described on this reference page is aligned with the ISO C standard. Any  
24436 conflict between the requirements described here and the ISO C standard is unintentional. This  
24437 volume of IEEE Std 1003.1-200x defers to the ISO C standard.24438 The *mbstowcs()* function shall convert a sequence of characters that begins in the initial shift  
24439 state from the array pointed to by *s* into a sequence of corresponding wide-character codes and  
24440 shall store not more than *n* wide-character codes into the array pointed to by *pwcs*. No  
24441 characters that follow a null byte (which is converted into a wide-character code with value 0)  
24442 shall be examined or converted. Each character shall be converted as if by a call to *mbtowc()*,  
24443 except that the shift state of *mbtowc()* is not affected.24444 No more than *n* elements shall be modified in the array pointed to by *pwcs*. If copying takes  
24445 place between objects that overlap, the behavior is undefined.24446 **XSI** The behavior of this function shall be affected by the *LC\_CTYPE* category of the current locale. If  
24447 *pwcs* is a null pointer, *mbstowcs()* shall return the length required to convert the entire array  
24448 regardless of the value of *n*, but no values are stored.24449 **RETURN VALUE**24450 **CX** If an invalid character is encountered, *mbstowcs()* shall return **(size\_t)-1** and may set *errno* to  
24451 **XSI** indicate the error. Otherwise, *mbstowcs()* shall return the number of the array elements modified  
24452 (or required if *pwcs* is null), not including a terminating 0 code, if any. The array shall not be  
24453 zero-terminated if the value returned is *n*.24454 **ERRORS**24455 The *mbstowcs()* function may fail if:24456 **XSI** [EILSEQ] Invalid byte sequence is detected.24457 **EXAMPLES**

24458 None.

24459 **APPLICATION USAGE**

24460 None.

24461 **RATIONALE**

24462 None.

24463 **FUTURE DIRECTIONS**

24464 None.

24465 **SEE ALSO**24466 *mblen()*, *mbtowc()*, *wctomb()*, *wcstombs()*, the Base Definitions volume of IEEE Std 1003.1-200x,  
24467 <stdlib.h>24468 **CHANGE HISTORY**

24469 First released in Issue 4. Aligned with the ISO C standard.

24470 **Issue 6**

24471 The *mbstowcs()* prototype is updated for alignment with the ISO/IEC 9899:1999 standard. |

24472 Extensions beyond the ISO C standard are now marked. |

## 24473 NAME

24474 mbtowc — convert a character to a wide-character code

## 24475 SYNOPSIS

24476 #include &lt;stdlib.h&gt;

24477 int mbtowc(wchar\_t \*restrict *pwc*, const char \*restrict *s*, size\_t *n*);

## 24478 DESCRIPTION

24479 CX The functionality described on this reference page is aligned with the ISO C standard. Any  
 24480 conflict between the requirements described here and the ISO C standard is unintentional. This  
 24481 volume of IEEE Std 1003.1-200x defers to the ISO C standard.

24482 If *s* is not a null pointer, *mbtowc()* shall determine the number of the bytes that constitute the  
 24483 character pointed to by *s*. It shall then determine the wide-character code for the value of type  
 24484 **wchar\_t** that corresponds to that character. (The value of the wide-character code corresponding  
 24485 to the null byte is 0.) If the character is valid and *pwc* is not a null pointer, *mbtowc()* shall store  
 24486 the wide-character code in the object pointed to by *pwc*.

24487 The behavior of this function is affected by the *LC\_CTYPE* category of the current locale. For a  
 24488 state-dependent encoding, this function is placed into its initial state by a call for which its  
 24489 character pointer argument, *s*, is a null pointer. Subsequent calls with *s* as other than a null  
 24490 pointer shall cause the internal state of the function to be altered as necessary. A call with *s* as a  
 24491 null pointer shall cause this function to return a non-zero value if encodings have state  
 24492 dependency, and 0 otherwise. If the implementation employs special bytes to change the shift  
 24493 state, these bytes shall not produce separate wide-character codes, but shall be grouped with an  
 24494 adjacent character. Changing the *LC\_CTYPE* category causes the shift state of this function to be  
 24495 unspecified. At most *n* bytes of the array pointed to by *s* shall be examined.

24496 The implementation shall behave as if no function defined in this volume of  
 24497 IEEE Std 1003.1-200x calls *mbtowc()*.

## 24498 RETURN VALUE

24499 If *s* is a null pointer, *mbtowc()* shall return a non-zero or 0 value, if character encodings,  
 24500 respectively, do or do not have state-dependent encodings. If *s* is not a null pointer, *mbtowc()*  
 24501 shall either return 0 (if *s* points to the null byte), or return the number of bytes that constitute the  
 24502 CX converted character (if the next *n* or fewer bytes form a valid character), or return -1 and may  
 24503 set *errno* to indicate the error (if they do not form a valid character).

24504 In no case shall the value returned be greater than *n* or the value of the {MB\_CUR\_MAX} macro.

## 24505 ERRORS

24506 The *mbtowc()* function may fail if:

24507 XSI [EILSEQ] Invalid character sequence is detected.

## 24508 EXAMPLES

24509 None.

## 24510 APPLICATION USAGE

24511 None.

## 24512 RATIONALE

24513 None.

## 24514 FUTURE DIRECTIONS

24515 None.

24516 **SEE ALSO**

24517            *mblen()*, *mbstowcs()*, *wctomb()*, *wcstombs()*, the Base Definitions volume of IEEE Std 1003.1-200x,  
24518            <**stdlib.h**>

24519 **CHANGE HISTORY**

24520            First released in Issue 4. Aligned with the ISO C standard.

24521 **Issue 6**

24522            The *mbtowc()* prototype is updated for alignment with the ISO/IEC 9899:1999 standard. |

24523            Extensions beyond the ISO C standard are now marked. |

24524 **NAME**

24525 memccpy — copy bytes in memory

24526 **SYNOPSIS**24527 XSI 

```
#include <string.h>
```

24528 

```
void *memccpy(void *restrict s1, const void *restrict s2,
```

  
24529 

```
int c, size_t n);
```

24530

24531 **DESCRIPTION**24532 The *memccpy()* function shall copy bytes from memory area *s2* into *s1*, stopping after the first  
24533 occurrence of byte *c* (converted to an **unsigned char**) is copied, or after *n* bytes are copied,  
24534 whichever comes first. If copying takes place between objects that overlap, the behavior is  
24535 undefined.24536 **RETURN VALUE**24537 The *memccpy()* function shall return a pointer to the byte after the copy of *c* in *s1*, or a null  
24538 pointer if *c* was not found in the first *n* bytes of *s2*.24539 **ERRORS**

24540 No errors are defined.

24541 **EXAMPLES**

24542 None.

24543 **APPLICATION USAGE**24544 The *memccpy()* function does not check for the overflow of the receiving memory area.24545 **RATIONALE**

24546 None.

24547 **FUTURE DIRECTIONS**

24548 None.

24549 **SEE ALSO**24550 The Base Definitions volume of IEEE Std 1003.1-200x, **<string.h>**24551 **CHANGE HISTORY**

24552 First released in Issue 1. Derived from Issue 1 of the SVID.

24553 **Issue 6**24554 The **restrict** keyword is added to the *memccpy()* prototype for alignment with the  
24555 ISO/IEC 9899:1999 standard.



24556 **NAME**

24557 memchr — find byte in memory

24558 **SYNOPSIS**

24559 #include &lt;string.h&gt;

24560 void \*memchr(const void \*s, int c, size\_t n);

24561 **DESCRIPTION**

24562 cx The functionality described on this reference page is aligned with the ISO C standard. Any  
24563 conflict between the requirements described here and the ISO C standard is unintentional. This  
24564 volume of IEEE Std 1003.1-200x defers to the ISO C standard.

24565 The *memchr()* function shall locate the first occurrence of *c* (converted to an **unsigned char**) in  
24566 the initial *n* bytes (each interpreted as **unsigned char**) of the object pointed to by *s*.

24567 **RETURN VALUE**

24568 The *memchr()* function shall return a pointer to the located byte, or a null pointer if the byte does  
24569 not occur in the object.

24570 **ERRORS**

24571 No errors are defined.

24572 **EXAMPLES**

24573 None.

24574 **APPLICATION USAGE**

24575 None.

24576 **RATIONALE**

24577 None.

24578 **FUTURE DIRECTIONS**

24579 None.

24580 **SEE ALSO**24581 The Base Definitions volume of IEEE Std 1003.1-200x, <**string.h**>24582 **CHANGE HISTORY**

24583 First released in Issue 1. Derived from Issue 1 of the SVID.

24584 **NAME**

24585       memcmp — compare bytes in memory

24586 **SYNOPSIS**

24587       #include &lt;string.h&gt;

24588       int memcmp(const void \*s1, const void \*s2, size\_t n);

24589 **DESCRIPTION**

24590 **CX**       The functionality described on this reference page is aligned with the ISO C standard. Any  
24591       conflict between the requirements described here and the ISO C standard is unintentional. This  
24592       volume of IEEE Std 1003.1-200x defers to the ISO C standard.

24593       The *memcmp()* function shall compare the first *n* bytes (each interpreted as **unsigned char**) of the  
24594       object pointed to by *s1* to the first *n* bytes of the object pointed to by *s2*.

24595       The sign of a non-zero return value shall be determined by the sign of the difference between the  
24596       values of the first pair of bytes (both interpreted as type **unsigned char**) that differ in the objects  
24597       being compared.

24598 **RETURN VALUE**

24599       The *memcmp()* function shall return an integer greater than, equal to, or less than 0, if the object  
24600       pointed to by *s1* is greater than, equal to, or less than the object pointed to by *s2*, respectively.

24601 **ERRORS**

24602       No errors are defined.

24603 **EXAMPLES**

24604       None.

24605 **APPLICATION USAGE**

24606       None.

24607 **RATIONALE**

24608       None.

24609 **FUTURE DIRECTIONS**

24610       None.

24611 **SEE ALSO**24612       The Base Definitions volume of IEEE Std 1003.1-200x, <**string.h**>24613 **CHANGE HISTORY**

24614       First released in Issue 1. Derived from Issue 1 of the SVID.

24615 **NAME**

24616           memcpy — copy bytes in memory

24617 **SYNOPSIS**

24618           #include <string.h>

24619           void \*memcpy(void \*restrict s1, const void \*restrict s2, size\_t n);

24620 **DESCRIPTION**

24621 **CX**       The functionality described on this reference page is aligned with the ISO C standard. Any  
24622           conflict between the requirements described here and the ISO C standard is unintentional. This  
24623           volume of IEEE Std 1003.1-200x defers to the ISO C standard.

24624           The *memcpy()* function shall copy *n* bytes from the object pointed to by *s2* into the object pointed  
24625           to by *s1*. If copying takes place between objects that overlap, the behavior is undefined.

24626 **RETURN VALUE**

24627           The *memcpy()* function shall return *s1*; no return value is reserved to indicate an error.

24628 **ERRORS**

24629           No errors are defined.

24630 **EXAMPLES**

24631           None.

24632 **APPLICATION USAGE**

24633           The *memcpy()* function does not check for the overflowing of the receiving memory area.

24634 **RATIONALE**

24635           None.

24636 **FUTURE DIRECTIONS**

24637           None.

24638 **SEE ALSO**

24639           The Base Definitions volume of IEEE Std 1003.1-200x, <**string.h**>

24640 **CHANGE HISTORY**

24641           First released in Issue 1. Derived from Issue 1 of the SVID.

24642 **Issue 6**

24643           The *memcpy()* prototype is updated for alignment with the ISO/IEC 9899:1999 standard.

24644 **NAME**

24645 memmove — copy bytes in memory with overlapping areas

24646 **SYNOPSIS**

24647 #include &lt;string.h&gt;

24648 void \*memmove(void \*s1, const void \*s2, size\_t n);

24649 **DESCRIPTION**

24650 **CX** The functionality described on this reference page is aligned with the ISO C standard. Any  
24651 conflict between the requirements described here and the ISO C standard is unintentional. This  
24652 volume of IEEE Std 1003.1-200x defers to the ISO C standard.

24653 The *memmove()* function shall copy *n* bytes from the object pointed to by *s2* into the object  
24654 pointed to by *s1*. Copying takes place as if the *n* bytes from the object pointed to by *s2* are first  
24655 copied into a temporary array of *n* bytes that does not overlap the objects pointed to by *s1* and  
24656 *s2*, and then the *n* bytes from the temporary array are copied into the object pointed to by *s1*.

24657 **RETURN VALUE**24658 The *memmove()* function shall return *s1*; no return value is reserved to indicate an error.24659 **ERRORS**

24660 No errors are defined.

24661 **EXAMPLES**

24662 None.

24663 **APPLICATION USAGE**

24664 None.

24665 **RATIONALE**

24666 None.

24667 **FUTURE DIRECTIONS**

24668 None.

24669 **SEE ALSO**24670 The Base Definitions volume of IEEE Std 1003.1-200x, <**string.h**>24671 **CHANGE HISTORY**

24672 First released in Issue 4. Derived from the ANSI C standard.

24673 **NAME**

24674           memset — set bytes in memory

24675 **SYNOPSIS**

24676           #include &lt;string.h&gt;

24677           void \*memset(void \*s, int c, size\_t n);

24678 **DESCRIPTION**

24679 cx       The functionality described on this reference page is aligned with the ISO C standard. Any  
24680       conflict between the requirements described here and the ISO C standard is unintentional. This  
24681       volume of IEEE Std 1003.1-200x defers to the ISO C standard.

24682       The *memset()* function shall copy *c* (converted to an **unsigned char**) into each of the first *n* bytes  
24683       of the object pointed to by *s*.

24684 **RETURN VALUE**24685       The *memset()* function shall return *s*; no return value is reserved to indicate an error.24686 **ERRORS**

24687       No errors are defined.

24688 **EXAMPLES**

24689       None.

24690 **APPLICATION USAGE**

24691       None.

24692 **RATIONALE**

24693       None.

24694 **FUTURE DIRECTIONS**

24695       None.

24696 **SEE ALSO**24697       The Base Definitions volume of IEEE Std 1003.1-200x, <**string.h**>24698 **CHANGE HISTORY**

24699       First released in Issue 1. Derived from Issue 1 of the SVID.

24700 **NAME**

24701 mkdir — make a directory

24702 **SYNOPSIS**

24703 #include <sys/stat.h>

24704 int mkdir(const char \*path, mode\_t mode);

24705 **DESCRIPTION**

24706 The *mkdir()* function shall create a new directory with name *path*. The file permission bits of the  
 24707 new directory shall be initialized from *mode*. These file permission bits of the *mode* argument  
 24708 shall be modified by the process' file creation mask.

24709 When bits in *mode* other than the file permission bits are set, the meaning of these additional bits  
 24710 is implementation-defined.

24711 The directory's user ID shall be set to the process' effective user ID. The directory's group ID  
 24712 shall be set to the group ID of the parent directory or to the effective group ID of the process.  
 24713 Implementations shall provide a way to initialize the directory's group ID to the group ID of the  
 24714 parent directory. Implementations may, but need not, provide an implementation-defined way  
 24715 to initialize the directory's group ID to the effective group ID of the calling process.

24716 The newly created directory shall be an empty directory.

24717 If *path* names a symbolic link, *mkdir()* shall fail and set *errno* to [EEXIST].

24718 Upon successful completion, *mkdir()* shall mark for update the *st\_atime*, *st\_ctime*, and *st\_mtime*  
 24719 fields of the directory. Also, the *st\_ctime* and *st\_mtime* fields of the directory that contains the  
 24720 new entry shall be marked for update.

24721 **RETURN VALUE**

24722 Upon successful completion, *mkdir()* shall return 0. Otherwise, -1 shall be returned, no directory  
 24723 shall be created, and *errno* shall be set to indicate the error.

24724 **ERRORS**

24725 The *mkdir()* function shall fail if:

24726 [EACCES] Search permission is denied on a component of the path prefix, or write  
 24727 permission is denied on the parent directory of the directory to be created.

24728 [EEXIST] The named file exists.

24729 [ELOOP] A loop exists in symbolic links encountered during resolution of the *path*  
 24730 argument.

24731 [EMLINK] The link count of the parent directory would exceed {LINK\_MAX}.

24732 [ENAMETOOLONG]

24733 The length of the *path* argument exceeds {PATH\_MAX} or a pathname  
 24734 component is longer than {NAME\_MAX}.

24735 [ENOENT] A component of the path prefix specified by *path* does not name an existing  
 24736 directory or *path* is an empty string.

24737 [ENOSPC] The file system does not contain enough space to hold the contents of the new  
 24738 directory or to extend the parent directory of the new directory.

24739 [ENOTDIR] A component of the path prefix is not a directory.

24740 [EROFS] The parent directory resides on a read-only file system.

24741 The *mkdir()* function may fail if:

24742 [ELOOP] More than {SYMLOOP\_MAX} symbolic links were encountered during  
24743 resolution of the *path* argument.

24744 [ENAMETOOLONG]

24745 As a result of encountering a symbolic link in resolution of the *path* argument, |  
24746 the length of the substituted pathname string exceeded {PATH\_MAX}. |

#### 24747 EXAMPLES

##### 24748 **Creating a Directory**

24749 The following example shows how to create a directory named */home/cnd/mod1*, with  
24750 read/write/search permissions for owner and group, and with read/search permissions for  
24751 others.

```
24752 #include <sys/types.h>
```

```
24753 #include <sys/stat.h>
```

```
24754 int status;
```

```
24755 ...
```

```
24756 status = mkdir("/home/cnd/mod1", S_IRWXU | S_IRWXG | S_IROTH | S_IXOTH);
```

#### 24757 APPLICATION USAGE

24758 None.

#### 24759 RATIONALE

24760 The *mkdir()* function originated in 4.2 BSD and was added to System V in Release 3.0.

24761 4.3 BSD detects [ENAMETOOLONG].

24762 The POSIX.1-1990 standard required that the group ID of a newly created directory be set to the |  
24763 group ID of its parent directory or to the effective group ID of the creating process. FIPS 151-2 |  
24764 required that implementations provide a way to have the group ID be set to the group ID of the |  
24765 containing directory, but did not prohibit implementations also supporting a way to set the |  
24766 group ID to the effective group ID of the creating process. Conforming applications should not |  
24767 assume which group ID will be used. If it matters, an application can use *chown()* to set the |  
24768 group ID after the directory is created, or determine under what conditions the implementation |  
24769 will set the desired group ID. |

#### 24770 FUTURE DIRECTIONS

24771 None.

#### 24772 SEE ALSO

24773 *umask()*, the Base Definitions volume of IEEE Std 1003.1-200x, *<sys/stat.h>*, *<sys/types.h>*

#### 24774 CHANGE HISTORY

24775 First released in Issue 3.

24776 Entry included for alignment with the POSIX.1-1988 standard.

#### 24777 Issue 6

24778 In the SYNOPSIS, the optional include of the *<sys/types.h>* header is removed.

24779 The following new requirements on POSIX implementations derive from alignment with the |  
24780 Single UNIX Specification:

- 24781 • The requirement to include *<sys/types.h>* has been removed. Although *<sys/types.h>* was  
24782 required for conforming implementations of previous POSIX specifications, it was not  
24783 required for UNIX applications.

24784

- The [ELOOP] mandatory error condition is added.

24785

- A second [ENAMETOOLONG] is added as an optional error condition.

24786

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

24787

- The [ELOOP] optional error condition is added.



24788 **NAME**

24789 mkfifo — make a FIFO special file

24790 **SYNOPSIS**

24791 #include &lt;sys/stat.h&gt;

24792 int mkfifo(const char \*path, mode\_t mode);

24793 **DESCRIPTION**

24794 The *mkfifo()* function shall create a new FIFO special file named by the pathname pointed to by |  
 24795 *path*. The file permission bits of the new FIFO shall be initialized from *mode*. The file permission |  
 24796 bits of the *mode* argument shall be modified by the process' file creation mask. |

24797 When bits in *mode* other than the file permission bits are set, the effect is implementation- |  
 24798 defined.

24799 If *path* names a symbolic link, *mkfifo()* shall fail and set *errno* to [EEXIST].

24800 The FIFO's user ID shall be set to the process' effective user ID. The FIFO's group ID shall be set |  
 24801 to the group ID of the parent directory or to the effective group ID of the process. |  
 24802 Implementation shall provide a way to initialize the FIFO's group ID to the group ID of the |  
 24803 parent directory. Implementations may, but need not, provide an implementation-defined way |  
 24804 to initialize the FIFO's group ID to the effective group ID of the calling process. |

24805 Upon successful completion, *mkfifo()* shall mark for update the *st\_atime*, *st\_ctime*, and *st\_mtime* |  
 24806 fields of the file. Also, the *st\_ctime* and *st\_mtime* fields of the directory that contains the new |  
 24807 entry shall be marked for update.

24808 **RETURN VALUE**

24809 Upon successful completion, 0 shall be returned. Otherwise, -1 shall be returned, no FIFO shall |  
 24810 be created, and *errno* shall be set to indicate the error.

24811 **ERRORS**24812 The *mkfifo()* function shall fail if:

24813 [EACCES] A component of the path prefix denies search permission, or write permission |  
 24814 is denied on the parent directory of the FIFO to be created.

24815 [EEXIST] The named file already exists.

24816 [ELOOP] A loop exists in symbolic links encountered during resolution of the *path* |  
 24817 argument.

24818 [ENAMETOOLONG]

24819 The length of the *path* argument exceeds {PATH\_MAX} or a pathname |  
 24820 component is longer than {NAME\_MAX}. |

24821 [ENOENT] A component of the path prefix specified by *path* does not name an existing |  
 24822 directory or *path* is an empty string.

24823 [ENOSPC] The directory that would contain the new file cannot be extended or the file |  
 24824 system is out of file-allocation resources.

24825 [ENOTDIR] A component of the path prefix is not a directory.

24826 [EROFS] The named file resides on a read-only file system.

24827 The *mkfifo()* function may fail if:

24828 [ELOOP] More than {SYMLOOP\_MAX} symbolic links were encountered during |  
 24829 resolution of the *path* argument.

24830 [ENAMETOOLONG]  
 24831 As a result of encountering a symbolic link in resolution of the *path* argument, |  
 24832 the length of the substituted pathname string exceeded {PATH\_MAX}. |

#### 24833 EXAMPLES

##### 24834 Creating a FIFO File

24835 The following example shows how to create a FIFO file named `/home/cnd/mod_done`, with  
 24836 read/write permissions for owner, and with read permissions for group and others.

```
24837 #include <sys/types.h>
24838 #include <sys/stat.h>
24839
24839 int status;
24840 ...
24841 status = mkfifo("/home/cnd/mod_done", S_IWUSR | S_IRUSR |
24842               S_IRGRP | S_IROTH);
```

#### 24843 APPLICATION USAGE

24844 None.

#### 24845 RATIONALE

24846 The syntax of this function is intended to maintain compatibility with historical  
 24847 implementations of *mknod()*. The latter function was included in the 1984 `/usr/group` standard  
 24848 but only for use in creating FIFO special files. The *mknod()* function was originally excluded  
 24849 from the POSIX.1-1988 standard as implementation-defined and replaced by *mkdir()* and  
 24850 *mkfifo()*. The *mknod()* function is now included for alignment with the Single UNIX  
 24851 Specification.

24852 The POSIX.1-1990 standard required that the group ID of a newly created FIFO be set to the |  
 24853 group ID of its parent directory or to the effective group ID of the creating process. FIPS 151-2 |  
 24854 required that implementations provide a way to have the group ID be set to the group ID of the |  
 24855 containing directory, but did not prohibit implementations also supporting a way to set the |  
 24856 group ID to the effective group ID of the creating process. Conforming applications should not |  
 24857 assume which group ID will be used. If it matters, an application can use *chown()* to set the |  
 24858 group ID after the FIFO is created, or determine under what conditions the implementation will |  
 24859 set the desired group ID. |

#### 24860 FUTURE DIRECTIONS

24861 None.

#### 24862 SEE ALSO

24863 *umask()*, the Base Definitions volume of IEEE Std 1003.1-200x, `<sys/stat.h>`, `<sys/types.h>`

#### 24864 CHANGE HISTORY

24865 First released in Issue 3.

24866 Entry included for alignment with the POSIX.1-1988 standard.

#### 24867 Issue 6

24868 In the SYNOPSIS, the optional include of the `<sys/types.h>` header is removed.

24869 The following new requirements on POSIX implementations derive from alignment with the |  
 24870 Single UNIX Specification:

- 24871 • The requirement to include `<sys/types.h>` has been removed. Although `<sys/types.h>` was  
 24872 required for conforming implementations of previous POSIX specifications, it was not  
 24873 required for UNIX applications.

- 24874           • The [ELOOP] mandatory error condition is added.
- 24875           • A second [ENAMETOOLONG] is added as an optional error condition.
- 24876       The following changes were made to align with the IEEE P1003.1a draft standard:
- 24877           • The [ELOOP] optional error condition is added.

24878 **NAME**

24879 **mknod** — make a directory, a special or regular file

24880 **SYNOPSIS**

24881 xSI `#include <sys/stat.h>`

24882 `int mknod(const char *path, mode_t mode, dev_t dev);`

24883

24884 **DESCRIPTION**

24885 The *mknod()* function shall create a new file named by the pathname to which the argument *path* |  
 24886 points.

24887 The file type for *path* is OR'ed into the *mode* argument, and the application shall select one of the  
 24888 following symbolic constants:

24889

24890

24891

24892

24893

24894

24895

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)

24896 The only portable use of *mknod()* is to create a FIFO-special file. If *mode* is not S\_IFIFO or *dev* is  
 24897 not 0, the behavior of *mknod()* is unspecified.

24898 The permissions for the new file are OR'ed into the *mode* argument, and may be selected from  
 24899 any combination of the following symbolic constants:

24900

24901

24902

24903

24904

24905

24906

24907

24908

24909

24910

24911

24912

24913

24914

24915

24916

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.

24917 The user ID of the file shall be initialized to the effective user ID of the process. The group ID of |  
 24918 the file shall be initialized to either the effective group ID of the process or the group ID of the |  
 24919 parent directory. Implementations shall provide a way to initialize the file's group ID to the |  
 24920 group ID of the parent directory. Implementations may, but need not, provide an |  
 24921 implementation-defined way to initialize the file's gorup ID to the effective group ID of the |  
 24922 calling proces. |

- 24923 The owner, group, and other permission bits of *mode* shall be modified by the file mode creation  
 24924 mask of the process. The *mknod()* function shall clear each bit whose corresponding bit in the file  
 24925 mode creation mask of the process is set.
- 24926 If *path* names a symbolic link, *mknod()* shall fail and set *errno* to [EEXIST].
- 24927 Upon successful completion, *mknod()* shall mark for update the *st\_atime*, *st\_ctime*, and *st\_mtime*  
 24928 fields of the file. Also, the *st\_ctime* and *st\_mtime* fields of the directory that contains the new  
 24929 entry shall be marked for update.
- 24930 Only a process with appropriate privileges may invoke *mknod()* for file types other than FIFO-  
 24931 special.
- 24932 **RETURN VALUE**
- 24933 Upon successful completion, *mknod()* shall return 0. Otherwise, it shall return -1, the new file  
 24934 shall not be created, and *errno* shall be set to indicate the error.
- 24935 **ERRORS**
- 24936 The *mknod()* function shall fail if:
- 24937 [EACCES] A component of the path prefix denies search permission, or write permission  
 24938 is denied on the parent directory.
- 24939 [EEXIST] The named file exists.
- 24940 [EINVAL] An invalid argument exists.
- 24941 [EIO] An I/O error occurred while accessing the file system.
- 24942 [ELOOP] A loop exists in symbolic links encountered during resolution of the *path*  
 24943 argument.
- 24944 [ENAMETOOLONG]  
 24945 The length of a pathname exceeds {PATH\_MAX} or a pathname component is  
 24946 longer than {NAME\_MAX}.
- 24947 [ENOENT] A component of the path prefix specified by *path* does not name an existing  
 24948 directory or *path* is an empty string.
- 24949 [ENOSPC] The directory that would contain the new file cannot be extended or the file  
 24950 system is out of file allocation resources.
- 24951 [ENOTDIR] A component of the path prefix is not a directory.
- 24952 [EPERM] The invoking process does not have appropriate privileges and the file type is  
 24953 not FIFO-special.
- 24954 [EROFS] The directory in which the file is to be created is located on a read-only file  
 24955 system.
- 24956 The *mknod()* function may fail if:
- 24957 [ELOOP] More than {SYMLOOP\_MAX} symbolic links were encountered during  
 24958 resolution of the *path* argument.
- 24959 [ENAMETOOLONG]  
 24960 Pathname resolution of a symbolic link produced an intermediate result  
 24961 whose length exceeds {PATH\_MAX}.

24962 **EXAMPLES**24963 **Creating a FIFO Special File**

24964 The following example shows how to create a FIFO special file named `/home/cnd/mod_done`,  
 24965 with read/write permissions for owner, and with read permissions for group and others.

```
24966 #include <sys/types.h>
24967 #include <sys/stat.h>

24968 dev_t dev;
24969 int status;
24970 ...
24971 status = mknod("/home/cnd/mod_done", S_IFIFO | S_IWUSR |
24972 S_IRUSR | S_IRGRP | S_IROTH, dev);
```

24973 **APPLICATION USAGE**

24974 `mkfifo()` is preferred over this function for making FIFO special files.

24975 **RATIONALE**

24976 The POSIX.1-1990 standard required that the group ID of a newly created file be set to the group  
 24977 ID of its parent directory or to the effective group ID of the creating process. FIPS 151-2 required  
 24978 that implementations provide a way to have the group ID be set to the group ID of the  
 24979 containing directory, but did not prohibit implementations also supporting a way to set the  
 24980 group ID to the effective group ID of the creating process. Conforming applications should not  
 24981 assume which group ID will be used. If it matters, an application can use `chown()` to set the  
 24982 group ID after the file is created, or determine under what conditions the implementation will  
 24983 set the desired group ID.

24984 **FUTURE DIRECTIONS**

24985 None.

24986 **SEE ALSO**

24987 `chmod()`, `creat()`, `exec`, `mkdir()`, `mkfifo()`, `open()`, `stat()`, `umask()`, the Base Definitions volume of  
 24988 IEEE Std 1003.1-200x, `<sys/stat.h>`

24989 **CHANGE HISTORY**

24990 First released in Issue 4, Version 2.

24991 **Issue 5**

24992 Moved from X/OPEN UNIX extension to BASE.

24993 **Issue 6**

24994 The DESCRIPTION is updated to avoid use of the term “must” for application requirements.

24995 The wording of the mandatory [ELOOP] error condition is updated, and a second optional  
 24996 [ELOOP] error condition is added.

24997 **NAME**

24998 mkstemp — make a unique filename

24999 **SYNOPSIS**25000 XSI `#include <stdlib.h>`25001 `int mkstemp(char *template);`

25002

25003 **DESCRIPTION**

25004 The *mkstemp()* function shall replace the contents of the string pointed to by *template* by a unique  
 25005 filename, and return a file descriptor for the file open for reading and writing. The function thus  
 25006 prevents any possible race condition between testing whether the file exists and opening it for  
 25007 use. The string in *template* should look like a filename with six trailing *Xs*; *mkstemp()* replaces  
 25008 each *X* with a character from the portable filename character set. The characters are chosen such  
 25009 that the resulting name does not duplicate the name of an existing file at the time of a call to  
 25010 *mkstemp()*.

25011 **RETURN VALUE**

25012 Upon successful completion, *mkstemp()* shall return an open file descriptor. Otherwise,  $-1$  shall  
 25013 be returned if no suitable file could be created.

25014 **ERRORS**

25015 No errors are defined.

25016 **EXAMPLES**25017 **Generating a Filename**

25018 The following example creates a file with a 10-character name beginning with the characters  
 25019 "file" and opens the file for reading and writing. The value returned as the value of *fd* is a file  
 25020 descriptor that identifies the file.

25021 `#include <stdlib.h>`25022 `...`25023 `char template[] = "/tmp/fileXXXXXX";`25024 `int fd;`25025 `fd = mkstemp(template);`25026 **APPLICATION USAGE**

25027 It is possible to run out of letters.

25028 The *mkstemp()* function need not check to determine whether the filename part of *template*  
 25029 exceeds the maximum allowable filename length.

25030 **RATIONALE**

25031 None.

25032 **FUTURE DIRECTIONS**

25033 None.

25034 **SEE ALSO**

25035 *getpid()*, *open()*, *tmpfile()*, *tmpnam()*, the Base Definitions volume of IEEE Std 1003.1-200x,  
 25036 `<stdlib.h>`

25037 **CHANGE HISTORY**

25038           First released in Issue 4, Version 2.

25039 **Issue 5**

25040           Moved from X/OPEN UNIX extension to BASE.



25041 **NAME**25042 mktemp — make a unique filename (**LEGACY**)25043 **SYNOPSIS**25044 XSI `#include <stdlib.h>`25045 `char *mktemp(char *template);`

25046

25047 **DESCRIPTION**

25048 The *mktemp()* function shall replace the contents of the string pointed to by *template* by a unique  
25049 filename and return *template*. The application shall initialize *template* to be a filename with six  
25050 trailing *X*s; *mktemp()* shall replace each *X* with a single byte character from the portable filename  
25051 character set.

25052 **RETURN VALUE**

25053 The *mktemp()* function shall return the pointer *template*. If a unique name cannot be created,  
25054 *template* shall point to a null string.

25055 **ERRORS**

25056 No errors are defined.

25057 **EXAMPLES**25058 **Generating a Filename**

25059 The following example replaces the contents of the "template" string with a 10-character  
25060 filename beginning with the characters "file" and returns a pointer to the "template" string  
25061 that contains the new filename.

25062 `#include <stdlib.h>`25063 `...`25064 `char *template = "/tmp/fileXXXXXX";`25065 `char *ptr;`25066 `ptr = mktemp(template);`25067 **APPLICATION USAGE**

25068 Between the time a pathname is created and the file opened, it is possible for some other process  
25069 to create a file with the same name. The *mkstemp()* function avoids this problem and is preferred  
25070 over this function.

25071 **RATIONALE**

25072 None.

25073 **FUTURE DIRECTIONS**

25074 This function may be withdrawn in a future version.

25075 **SEE ALSO**25076 *mkstemp()*, *tmpfile()*, *tmpnam()*, the Base Definitions volume of IEEE Std 1003.1-200x, `<stdlib.h>`25077 **CHANGE HISTORY**

25078 First released in Issue 4, Version 2.

25079 **Issue 5**

25080 Moved from X/OPEN UNIX extension to BASE.

25081 **Issue 6**

25082 This function is marked LEGACY.

25083 The DESCRIPTION is updated to avoid use of the term “must” for application requirements.

25084 **NAME**

25085           mktime — convert broken-down time into time since the Epoch

25086 **SYNOPSIS**

25087           #include &lt;time.h&gt;

25088           time\_t mktime(struct tm \*timeptr);

25089 **DESCRIPTION**

25090 **CX**       The functionality described on this reference page is aligned with the ISO C standard. Any  
 25091 conflict between the requirements described here and the ISO C standard is unintentional. This  
 25092 volume of IEEE Std 1003.1-200x defers to the ISO C standard.

25093       The *mktime()* function shall convert the broken-down time, expressed as local time, in the  
 25094 structure pointed to by *timeptr*, into a time since the Epoch value with the same encoding as that  
 25095 of the values returned by *time()*. The original values of the *tm\_wday* and *tm\_yday* components of  
 25096 the structure are ignored, and the original values of the other components are not restricted to  
 25097 the ranges described in <**time.h**>.

25098 **CX**       A positive or 0 value for *tm\_isdst* shall cause *mktime()* to presume initially that Daylight Savings  
 25099 Time, respectively, is or is not in effect for the specified time. A negative value for *tm\_isdst* shall  
 25100 cause *mktime()* to attempt to determine whether Daylight Saving Time is in effect for the  
 25101 specified time.

25102       Local timezone information shall be set as though *mktime()* called *tzset()*.

25103       The relationship between the **tm** structure (defined in the <**time.h**> header) and the time in  
 25104 seconds since the Epoch is that the result shall be as specified in the expression given in the  
 25105 definition of seconds since the Epoch (see the Base Definitions volume of IEEE Std 1003.1-200x,  
 25106 Section 4.14, Seconds Since the Epoch) corrected for timezone and any seasonal time  
 25107 adjustments, where the names in the structure and in the expression correspond.

25108       Upon successful completion, the values of the *tm\_wday* and *tm\_yday* components of the structure  
 25109 shall be set appropriately, and the other components are set to represent the specified time since  
 25110 the Epoch, but with their values forced to the ranges indicated in the <**time.h**> entry; the final  
 25111 value of *tm\_mday* shall not be set until *tm\_mon* and *tm\_year* are determined.

25112 **RETURN VALUE**

25113       The *mktime()* function shall return the specified time since the Epoch encoded as a value of type  
 25114 **time\_t**. If the time since the Epoch cannot be represented, the function shall return the value  
 25115 (**time\_t**)-1.

25116 **ERRORS**

25117       No errors are defined.

25118 **EXAMPLES**

25119       What day of the week is July 4, 2001?

25120       #include &lt;stdio.h&gt;

25121       #include &lt;time.h&gt;

25122       struct tm time\_str;

25123       char daybuf[20];

25124       int main(void)

25125       {

25126           time\_str.tm\_year = 2001 - 1900;

25127           time\_str.tm\_mon = 7 - 1;

25128           time\_str.tm\_mday = 4;

```
25129         time_str.tm_hour = 0;
25130         time_str.tm_min = 0;
25131         time_str.tm_sec = 1;
25132         time_str.tm_isdst = -1;
25133         if (mktime(&time_str) == -1)
25134             (void)puts("-unknown-");
25135         else {
25136             (void)strftime(daybuf, sizeof(daybuf), "%A", &time_str);
25137             (void)puts(daybuf);
25138         }
25139         return 0;
25140     }
```

**25141 APPLICATION USAGE**

25142 None.

**25143 RATIONALE**

25144 None.

**25145 FUTURE DIRECTIONS**

25146 None.

**25147 SEE ALSO**

25148 *asctime()*, *clock()*, *ctime()*, *difftime()*, *gmtime()*, *localtime()*, *strftime()*, *strptime()*, *time()*, *utime()*,  
25149 the Base Definitions volume of IEEE Std 1003.1-200x, <**time.h**>

**25150 CHANGE HISTORY**

25151 First released in Issue 3.

25152 Entry included for alignment with the POSIX.1-1988 standard and the ANSI C standard.

**25153 Issue 6**

25154 Extensions beyond the ISO C standard are now marked.

25155 **NAME**25156 mlock, munlock — lock or unlock a range of process address space (**REALTIME**)25157 **SYNOPSIS**

25158 MLR #include &lt;sys/mman.h&gt;

25159 int mlock(const void \*addr, size\_t len);

25160 int munlock(const void \*addr, size\_t len);

25161

25162 **DESCRIPTION**

25163 The *mlock()* function shall cause those whole pages containing any part of the address space of  
 25164 the process starting at address *addr* and continuing for *len* bytes to be memory-resident until  
 25165 unlocked or until the process exits or *execs* another process image. The implementation may  
 25166 require that *addr* be a multiple of {PAGESIZE}.

25167 The *munlock()* function shall unlock those whole pages containing any part of the address space  
 25168 of the process starting at address *addr* and continuing for *len* bytes, regardless of how many  
 25169 times *mlock()* has been called by the process for any of the pages in the specified range. The  
 25170 implementation may require that *addr* be a multiple of {PAGESIZE}.

25171 If any of the pages in the range specified to a call to *munlock()* are also mapped into the address  
 25172 spaces of other processes, any locks established on those pages by another process are  
 25173 unaffected by the call of this process to *munlock()*. If any of the pages in the range specified by a  
 25174 call to *munlock()* are also mapped into other portions of the address space of the calling process  
 25175 outside the range specified, any locks established on those pages via the other mappings are also  
 25176 unaffected by this call.

25177 Upon successful return from *mlock()*, pages in the specified range shall be locked and memory-  
 25178 resident. Upon successful return from *munlock()*, pages in the specified range shall be unlocked  
 25179 with respect to the address space of the process. Memory residency of unlocked pages is  
 25180 unspecified.

25181 The appropriate privilege is required to lock process memory with *mlock()*.

25182 **RETURN VALUE**

25183 Upon successful completion, the *mlock()* and *munlock()* functions shall return a value of zero.  
 25184 Otherwise, no change is made to any locks in the address space of the process, and the function  
 25185 shall return a value of  $-1$  and set *errno* to indicate the error.

25186 **ERRORS**

25187 The *mlock()* and *munlock()* functions shall fail if:

25188 [ENOMEM] Some or all of the address range specified by the *addr* and *len* arguments does  
 25189 not correspond to valid mapped pages in the address space of the process.

25190 The *mlock()* function shall fail if:

25191 [EAGAIN] Some or all of the memory identified by the operation could not be locked  
 25192 when the call was made.

25193 The *mlock()* and *munlock()* functions may fail if:

25194 [EINVAL] The *addr* argument is not a multiple of {PAGESIZE}.

25195 The *mlock()* function may fail if:

25196 [ENOMEM] Locking the pages mapped by the specified range would exceed an  
 25197 implementation-defined limit on the amount of memory that the process may  
 25198 lock.

25199 [EPERM] The calling process does not have the appropriate privilege to perform the  
25200 requested operation.

25201 **EXAMPLES**

25202 None.

25203 **APPLICATION USAGE**

25204 None.

25205 **RATIONALE**

25206 None.

25207 **FUTURE DIRECTIONS**

25208 None.

25209 **SEE ALSO**

25210 *exec*, *exit()*, *fork()*, *mlockall()*, *munmap()*, the Base Definitions volume of IEEE Std 1003.1-200x,  
25211 <sys/mman.h>

25212 **CHANGE HISTORY**

25213 First released in Issue 5. Included for alignment with the POSIX Realtime Extension.

25214 **Issue 6**

25215 The *mlock()* and *munlock()* functions are marked as part of the Range Memory Locking option.

25216 The [ENOSYS] error condition has been removed as stubs need not be provided if an  
25217 implementation does not support the Range Memory Locking option.

25218 **NAME**25219 mlockall, munlockall — lock/unlock the address space of a process (**REALTIME**)25220 **SYNOPSIS**

```
25221 ML      #include <sys/mman.h>
25222          int mlockall(int flags);
25223          int munlockall(void);
25224
```

25225 **DESCRIPTION**

25226 The *mlockall()* function shall cause all of the pages mapped by the address space of a process to  
 25227 be memory-resident until unlocked or until the process exits or *execs* another process image. The  
 25228 *flags* argument determines whether the pages to be locked are those currently mapped by the  
 25229 address space of the process, those that are mapped in the future, or both. The *flags* argument is  
 25230 constructed from the bitwise-inclusive OR of one or more of the following symbolic constants,  
 25231 defined in *<sys/mman.h>*:

25232 **MCL\_CURRENT** Lock all of the pages currently mapped into the address space of the process.

25233 **MCL\_FUTURE** Lock all of the pages that become mapped into the address space of the  
 25234 process in the future, when those mappings are established.

25235 If **MCL\_FUTURE** is specified, and the automatic locking of future mappings eventually causes  
 25236 the amount of locked memory to exceed the amount of available physical memory or any other  
 25237 implementation-defined limit, the behavior is implementation-defined. The manner in which the  
 25238 implementation informs the application of these situations is also implementation-defined.

25239 The *munlockall()* function shall unlock all currently mapped pages of the address space of the  
 25240 process. Any pages that become mapped into the address space of the process after a call to  
 25241 *munlockall()* shall not be locked, unless there is an intervening call to *mlockall()* specifying  
 25242 **MCL\_FUTURE** or a subsequent call to *mlockall()* specifying **MCL\_CURRENT**. If pages mapped  
 25243 into the address space of the process are also mapped into the address spaces of other processes  
 25244 and are locked by those processes, the locks established by the other processes shall be  
 25245 unaffected by a call by this process to *munlockall()*.

25246 Upon successful return from the *mlockall()* function that specifies **MCL\_CURRENT**, all currently  
 25247 mapped pages of the process' address space shall be memory-resident and locked. Upon return  
 25248 from the *munlockall()* function, all currently mapped pages of the process' address space shall be  
 25249 unlocked with respect to the process' address space. The memory residency of unlocked pages is  
 25250 unspecified.

25251 The appropriate privilege is required to lock process memory with *mlockall()*.

25252 **RETURN VALUE**

25253 Upon successful completion, the *mlockall()* function shall return a value of zero. Otherwise, no  
 25254 additional memory shall be locked, and the function shall return a value of  $-1$  and set *errno* to  
 25255 indicate the error. The effect of failure of *mlockall()* on previously existing locks in the address  
 25256 space is unspecified.

25257 If it is supported by the implementation, the *munlockall()* function shall always return a value of  
 25258 zero. Otherwise, the function shall return a value of  $-1$  and set *errno* to indicate the error.

25259 **ERRORS**

25260 The *mlockall()* function shall fail if:

25261 [EAGAIN] Some or all of the memory identified by the operation could not be locked  
 25262 when the call was made.

- 25263 [EINVAL] The *flags* argument is zero, or includes unimplemented flags.
- 25264 The *mlockall()* function may fail if:
- 25265 [ENOMEM] Locking all of the pages currently mapped into the address space of the  
25266 process would exceed an implementation-defined limit on the amount of  
25267 memory that the process may lock.
- 25268 [EPERM] The calling process does not have the appropriate privilege to perform the  
25269 requested operation.
- 25270 **EXAMPLES**
- 25271 None.
- 25272 **APPLICATION USAGE**
- 25273 None.
- 25274 **RATIONALE**
- 25275 None.
- 25276 **FUTURE DIRECTIONS**
- 25277 None.
- 25278 **SEE ALSO**
- 25279 *exec*, *exit()*, *fork()*, *mlock()*, *munmap()*, the Base Definitions volume of IEEE Std 1003.1-200x,  
25280 <*sys/mman.h*>
- 25281 **CHANGE HISTORY**
- 25282 First released in Issue 5. Included for alignment with the POSIX Realtime Extension.
- 25283 **Issue 6**
- 25284 The *mlockall()* and *munlockall()* functions are marked as part of the Process Memory Locking  
25285 option.
- 25286 The [ENOSYS] error condition has been removed as stubs need not be provided if an  
25287 implementation does not support the Process Memory Locking option.



## 25288 NAME

25289 mmap — map pages of memory

## 25290 SYNOPSIS

25291 MF|SHM #include &lt;sys/mman.h&gt;

```
25292 void *mmap(void *addr, size_t len, int prot, int flags,
25293           int fildes, off_t off);
25294
```

## 25295 DESCRIPTION

25296 The *mmap()* function shall establish a mapping between a process' address space and a file, |  
 25297 TYM shared memory object, or typed memory object. The format of the call is as follows:

```
25298 pa=mmap(addr, len, prot, flags, fildes, off);
```

25299 The *mmap()* function shall establish a mapping between the address space of the process at an |  
 25300 address *pa* for *len* bytes to the memory object represented by the file descriptor *fildes* at offset *off* |  
 25301 for *len* bytes. The value of *pa* is an implementation-defined function of the parameter *addr* and |  
 25302 the values of *flags*, further described below. A successful *mmap()* call shall return *pa* as its result. |  
 25303 The address range starting at *pa* and continuing for *len* bytes shall be legitimate for the possible |  
 25304 (not necessarily current) address space of the process. The range of bytes starting at *off* and |  
 25305 continuing for *len* bytes shall be legitimate for the possible (not necessarily current) offsets in the |  
 25306 TYM file, shared memory object, or typed memory object represented by *fildes*.

25307 TYM If *fildes* represents a typed memory object opened with either the |  
 25308 POSIX\_TYPED\_MEM\_ALLOCATE flag or the POSIX\_TYPED\_MEM\_ALLOCATE\_CONTIG |  
 25309 flag, the memory object to be mapped shall be that portion of the typed memory object allocated |  
 25310 by the implementation as specified below. In this case, if *off* is non-zero, the behavior of *mmap()* |  
 25311 is undefined. If *fildes* refers to a valid typed memory object that is not accessible from the calling |  
 25312 process, *mmap()* shall fail.

25313 The mapping established by *mmap()* shall replace any previous mappings for those whole pages |  
 25314 containing any part of the address space of the process starting at *pa* and continuing for *len* |  
 25315 bytes.

25316 If the size of the mapped file changes after the call to *mmap()* as a result of some other operation |  
 25317 on the mapped file, the effect of references to portions of the mapped region that correspond to |  
 25318 added or removed portions of the file is unspecified.

25319 TYM The *mmap()* function shall be supported for regular files, shared memory objects, and typed |  
 25320 memory objects. Support for any other type of file is unspecified.

25321 The parameter *prot* determines whether read, write, execute, or some combination of accesses |  
 25322 are permitted to the data being mapped. The *prot* shall be either PROT\_NONE or the bitwise- |  
 25323 inclusive OR of one or more of the other flags in the following table, defined in the |  
 25324 <sys/mman.h> header.

25325

25326

25327

25328

25329

25330

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.

25331 If an implementation cannot support the combination of access types specified by *prot*, the call |  
 25332 MPR to *mmap()* shall fail. An implementation may permit accesses other than those specified by *prot*; |  
 25333 however, if the Memory Protection option is supported, the implementation shall not permit a

25334 write to succeed where PROT\_WRITE has not been set or shall not permit any access where  
 25335 PROT\_NONE alone has been set. The implementation shall support at least the following values  
 25336 of *prot*: PROT\_NONE, PROT\_READ, PROT\_WRITE, and the bitwise-inclusive OR of  
 25337 PROT\_READ and PROT\_WRITE. If the Memory Protection option is not supported, the result of  
 25338 any access that conflicts with the specified protection is undefined. The file descriptor *fildes* shall  
 25339 have been opened with read permission, regardless of the protection options specified. If  
 25340 PROT\_WRITE is specified, the application shall ensure that it has opened the file descriptor  
 25341 *fildes* with write permission unless MAP\_PRIVATE is specified in the *flags* parameter as  
 25342 described below.

25343 The parameter *flags* provides other information about the handling of the mapped data. The  
 25344 value of *flags* is the bitwise-inclusive OR of these options, defined in <sys/mman.h>:

25345

25346

25347

25348

25349

Symbolic Constant	Description
MAP_SHARED	Changes are shared.
MAP_PRIVATE	Changes are private.
MAP_FIXED	Interpret <i>addr</i> exactly.

25350 Implementations that do not support the Memory Mapped Files option are not required to  
 25351 support MAP\_PRIVATE.

25352 XSI It is implementation-defined whether MAP\_FIXED shall be supported. MAP\_FIXED shall be  
 25353 supported on XSI-conformant systems.

25354 MAP\_SHARED and MAP\_PRIVATE describe the disposition of write references to the memory  
 25355 object. If MAP\_SHARED is specified, write references shall change the underlying object. If  
 25356 MAP\_PRIVATE is specified, modifications to the mapped data by the calling process shall be  
 25357 visible only to the calling process and shall not change the underlying object. It is unspecified  
 25358 whether modifications to the underlying object done after the MAP\_PRIVATE mapping is  
 25359 established are visible through the MAP\_PRIVATE mapping. Either MAP\_SHARED or  
 25360 MAP\_PRIVATE can be specified, but not both. The mapping type is retained across *fork*().

25361 TYM When *fildes* represents a typed memory object opened with either the  
 25362 POSIX\_TYPED\_MEM\_ALLOCATE flag or the POSIX\_TYPED\_MEM\_ALLOCATE\_CONTIG  
 25363 flag, *mmap*() shall, if there are enough resources available, map *len* bytes allocated from the  
 25364 corresponding typed memory object which were not previously allocated to any process in any  
 25365 processor that may access that typed memory object. If there are not enough resources available,  
 25366 the function shall fail. If *fildes* represents a typed memory object opened with the  
 25367 POSIX\_TYPED\_MEM\_ALLOCATE\_CONTIG flag, these allocated bytes shall be contiguous  
 25368 within the typed memory object. If *fildes* represents a typed memory object opened with the  
 25369 POSIX\_TYPED\_MEM\_ALLOCATE flag, these allocated bytes may be composed of non-  
 25370 contiguous fragments within the typed memory object. If *fildes* represents a typed memory  
 25371 object opened with neither the POSIX\_TYPED\_MEM\_ALLOCATE\_CONTIG flag nor the  
 25372 POSIX\_TYPED\_MEM\_ALLOCATE flag, *len* bytes starting at offset *off* within the typed memory  
 25373 object are mapped, exactly as when mapping a file or shared memory object. In this case, if two  
 25374 processes map an area of typed memory using the same *off* and *len* values and using file  
 25375 descriptors that refer to the same memory pool (either from the same port or from a different  
 25376 port), both processes shall map the same region of storage.

25377 When MAP\_FIXED is set in the *flags* argument, the implementation is informed that the value of  
 25378 *pa* shall be *addr*, exactly. If MAP\_FIXED is set, *mmap*() may return MAP\_FAILED and set *errno* to  
 25379 [EINVAL]. If a MAP\_FIXED request is successful, the mapping established by *mmap*() replaces  
 25380 any previous mappings for the process' pages in the range [*pa*,*pa+len*).

25381 When MAP\_FIXED is not set, the implementation uses *addr* in an implementation-defined  
 25382 manner to arrive at *pa*. The *pa* so chosen shall be an area of the address space that the  
 25383 implementation deems suitable for a mapping of *len* bytes to the file. All implementations  
 25384 interpret an *addr* value of 0 as granting the implementation complete freedom in selecting *pa*,  
 25385 subject to constraints described below. A non-zero value of *addr* is taken to be a suggestion of a  
 25386 process address near which the mapping should be placed. When the implementation selects a  
 25387 value for *pa*, it never places a mapping at address 0, nor does it replace any extant mapping.

25388 The *off* argument is constrained to be aligned and sized according to the value returned by  
 25389 *sysconf()* when passed *\_SC\_PAGESIZE* or *\_SC\_PAGE\_SIZE*. When MAP\_FIXED is specified, the  
 25390 application shall ensure that the argument *addr* also meets these constraints. The  
 25391 implementation performs mapping operations over whole pages. Thus, while the argument *len*  
 25392 need not meet a size or alignment constraint, the implementation shall include, in any mapping  
 25393 operation, any partial page specified by the range [*pa*,*pa+len*).

25394 The system shall always zero-fill any partial page at the end of an object. Further, the system  
 25395 shall never write out any modified portions of the last page of an object which are beyond its  
 25396 MPR end. References within the address range starting at *pa* and continuing for *len* bytes to whole  
 25397 pages following the end of an object shall result in delivery of a SIGBUS signal.

25398 An implementation may generate SIGBUS signals when a reference would cause an error in the  
 25399 mapped object, such as out-of-space condition.

25400 The *mmap()* function shall add an extra reference to the file associated with the file descriptor  
 25401 *fdes* which is not removed by a subsequent *close()* on that file descriptor. This reference shall be  
 25402 removed when there are no more mappings to the file. |

25403 The *st\_atime* field of the mapped file may be marked for update at any time between the *mmap()*  
 25404 call and the corresponding *munmap()* call. The initial read or write reference to a mapped region  
 25405 shall cause the file's *st\_atime* field to be marked for update if it has not already been marked for  
 25406 update.

25407 The *st\_ctime* and *st\_mtime* fields of a file that is mapped with MAP\_SHARED and PROT\_WRITE  
 25408 shall be marked for update at some point in the interval between a write reference to the  
 25409 mapped region and the next call to *msync()* with MS\_ASYNC or MS\_SYNC for that portion of  
 25410 the file by any process. If there is no such call and if the underlying file is modified as a result of  
 25411 a write reference, then these fields shall be marked for update at some time after the write  
 25412 reference.

25413 There may be implementation-defined limits on the number of memory regions that can be  
 25414 mapped (per process or per system).

25415 XSI If such a limit is imposed, whether the number of memory regions that can be mapped by a  
 25416 process is decreased by the use of *shmat()* is implementation-defined.

25417 If *mmap()* fails for reasons other than [EBADF], [EINVAL], or [ENOTSUP], some of the  
 25418 mappings in the address range starting at *addr* and continuing for *len* bytes may have been  
 25419 unmapped.

#### 25420 RETURN VALUE

25421 Upon successful completion, the *mmap()* function shall return the address at which the mapping  
 25422 was placed (*pa*); otherwise, it shall return a value of MAP\_FAILED and set *errno* to indicate the  
 25423 error. The symbol MAP\_FAILED is defined in the <sys/mman.h> header. No successful return  
 25424 from *mmap()* shall return the value MAP\_FAILED. |

25425 **ERRORS**

- 25426 The *mmap()* function shall fail if:
- 25427 [EACCES] The *fildev* argument is not open for read, regardless of the protection specified,  
25428 or *fildev* is not open for write and PROT\_WRITE was specified for a  
25429 MAP\_SHARED type mapping.
- 25430 ML [EAGAIN] The mapping could not be locked in memory, if required by *mlockall()*, due to  
25431 a lack of resources.
- 25432 [EBADF] The *fildev* argument is not a valid open file descriptor.
- 25433 [EINVAL] The *addr* argument (if MAP\_FIXED was specified) or *off* is not a multiple of  
25434 the page size as returned by *sysconf()*, or are considered invalid by the  
25435 implementation.
- 25436 [EINVAL] The value of *flags* is invalid (neither MAP\_PRIVATE nor MAP\_SHARED is  
25437 set).
- 25438 [EMFILE] The number of mapped regions would exceed an implementation-defined  
25439 limit (per process or per system).
- 25440 [ENODEV] The *fildev* argument refers to a file whose type is not supported by *mmap()*.
- 25441 [ENOMEM] MAP\_FIXED was specified, and the range [*addr,addr+len*) exceeds that allowed  
25442 for the address space of a process; or, if MAP\_FIXED was not specified and  
25443 there is insufficient room in the address space to effect the mapping.
- 25444 ML [ENOMEM] The mapping could not be locked in memory, if required by *mlockall()*,  
25445 because it would require more space than the system is able to supply.
- 25446 MAP\_FIXED or MAP\_PRIVATE was specified in the *flags* argument and the  
25447 implementation does not support this functionality.
- 25448 TYM [ENOMEM] Not enough unallocated memory resources remain in the typed memory  
25449 object designated by *fildev* to allocate *len* bytes.
- 25450 [ENOTSUP] The implementation does not support the combination of accesses requested  
25451 in the *prot* argument.
- 25452 [ENXIO] Addresses in the range [*off,off+len*) are invalid for the object specified by *fildev*.
- 25453 [ENXIO] MAP\_FIXED was specified in *flags* and the combination of *addr*, *len*, and *off* is  
25454 invalid for the object specified by *fildev*.
- 25455 TYM [ENXIO] The *fildev* argument refers to a typed memory object that is not accessible from  
25456 the calling process.
- 25457 [EOVERFLOW] The file is a regular file and the value of *off* plus *len* exceeds the offset  
25458 maximum established in the open file description associated with *fildev*.

25459 **EXAMPLES**

25460 None.

25461 **APPLICATION USAGE**

25462 Use of *mmap()* may reduce the amount of memory available to other memory allocation  
25463 functions.

25464 Use of MAP\_FIXED may result in unspecified behavior in further use of *malloc()* and *shmat()*.  
25465 The use of MAP\_FIXED is discouraged, as it may prevent an implementation from making the  
25466 most effective use of resources.

25467 The application must ensure correct synchronization when using *mmap()* in conjunction with  
 25468 any other file access method, such as *read()* and *write()*, standard input/output, and *shmat()*.

25469 The *mmap()* function allows access to resources via address space manipulations, instead of  
 25470 *read()/write()*. Once a file is mapped, all a process has to do to access it is use the data at the  
 25471 address to which the file was mapped. So, using pseudo-code to illustrate the way in which an  
 25472 existing program might be changed to use *mmap()*, the following:

```
25473 fildes = open(...)
25474 lseek(fildes, some_offset)
25475 read(fildes, buf, len)
25476 /* Use data in buf. */
```

25477 becomes:

```
25478 fildes = open(...)
25479 address = mmap(0, len, PROT_READ, MAP_PRIVATE, fildes, some_offset)
25480 /* Use data at address. */
```

25481 The [EINVAL] error above is marked EX because it is defined as an optional error in the POSIX  
 25482 Realtime Extension.

#### 25483 RATIONALE

25484 After considering several other alternatives, it was decided to adopt the *mmap()* definition found  
 25485 in SVR4 for mapping memory objects into process address spaces. The SVR4 definition is  
 25486 minimal, in that it describes only what has been built, and what appears to be necessary for a  
 25487 general and portable mapping facility.

25488 Note that while *mmap()* was first designed for mapping files, it is actually a general-purpose  
 25489 mapping facility. It can be used to map any appropriate object, such as memory, files, devices,  
 25490 and so on, into the address space of a process.

25491 When a mapping is established, it is possible that the implementation may need to map more  
 25492 than is requested into the address space of the process because of hardware requirements. An  
 25493 application, however, cannot count on this behavior. Implementations that do not use a paged  
 25494 architecture may simply allocate a common memory region and return the address of it; such  
 25495 implementations probably do not allocate any more than is necessary. References past the end of  
 25496 the requested area are unspecified.

25497 If an application requests a mapping that would overlay existing mappings in the process, it  
 25498 might be desirable that an implementation detect this and inform the application. However, the  
 25499 default, portable (not MAP\_FIXED) operation does not overlay existing mappings. On the other  
 25500 hand, if the program specifies a fixed address mapping (which requires some implementation  
 25501 knowledge to determine a suitable address, if the function is supported at all), then the program  
 25502 is presumed to be successfully managing its own address space and should be trusted when it  
 25503 asks to map over existing data structures. Furthermore, it is also desirable to make as few system  
 25504 calls as possible, and it might be considered onerous to require an *munmap()* before an *mmap()*  
 25505 to the same address range. This volume of IEEE Std 1003.1-200x specifies that the new mappings  
 25506 replace any existing mappings, following existing practice in this regard.

25507 It is not expected, when the Memory Protection option is supported, that all hardware  
 25508 implementations are able to support all combinations of permissions at all addresses. When this  
 25509 option is supported, implementations are required to disallow write access to mappings without  
 25510 write permission and to disallow access to mappings without any access permission. Other than  
 25511 these restrictions, implementations may allow access types other than those requested by the  
 25512 application. For example, if the application requests only PROT\_WRITE, the implementation  
 25513 may also allow read access. A call to *mmap()* fails if the implementation cannot support allowing

25514 all the access requested by the application. For example, some implementations cannot support  
25515 a request for both write access and execute access simultaneously. All implementations  
25516 supporting the Memory Protection option must support requests for no access, read access,  
25517 write access, and both read and write access. Strictly conforming code must only rely on the  
25518 required checks. These restrictions allow for portability across a wide range of hardware.

25519 The MAP\_FIXED address treatment is likely to fail for non-page-aligned values and for certain  
25520 architecture-dependent address ranges. Conforming implementations cannot count on being  
25521 able to choose address values for MAP\_FIXED without utilizing non-portable, implementation-  
25522 defined knowledge. Nonetheless, MAP\_FIXED is provided as a standard interface conforming to  
25523 existing practice for utilizing such knowledge when it is available.

25524 Similarly, in order to allow implementations that do not support virtual addresses, support for  
25525 directly specifying any mapping addresses via MAP\_FIXED is not required and thus a  
25526 conforming application may not count on it.

25527 The MAP\_PRIVATE function can be implemented efficiently when memory protection hardware  
25528 is available. When such hardware is not available, implementations can implement such  
25529 “mappings” by simply making a real copy of the relevant data into process private memory,  
25530 though this tends to behave similarly to *read()*.

25531 The function has been defined to allow for many different models of using shared memory.  
25532 However, all uses are not equally portable across all machine architectures. In particular, the  
25533 *mmap()* function allows the system as well as the application to specify the address at which to  
25534 map a specific region of a memory object. The most portable way to use the function is always to  
25535 let the system choose the address, specifying NULL as the value for the argument *addr* and not  
25536 to specify MAP\_FIXED.

25537 If it is intended that a particular region of a memory object be mapped at the same address in a  
25538 group of processes (on machines where this is even possible), then MAP\_FIXED can be used to  
25539 pass in the desired mapping address. The system can still be used to choose the desired address  
25540 if the first such mapping is made without specifying MAP\_FIXED, and then the resulting  
25541 mapping address can be passed to subsequent processes for them to pass in via MAP\_FIXED.  
25542 The availability of a specific address range cannot be guaranteed, in general.

25543 The *mmap()* function can be used to map a region of memory that is larger than the current size  
25544 of the object. Memory access within the mapping but beyond the current end of the underlying  
25545 objects may result in SIGBUS signals being sent to the process. The reason for this is that the size  
25546 of the object can be manipulated by other processes and can change at any moment. The  
25547 implementation should tell the application that a memory reference is outside the object where  
25548 this can be detected; otherwise, written data may be lost and read data may not reflect actual  
25549 data in the object.

25550 Note that references beyond the end of the object do not extend the object as the new end cannot  
25551 be determined precisely by most virtual memory hardware. Instead, the size can be directly  
25552 manipulated by *ftruncate()*.

25553 Process memory locking does apply to shared memory regions, and the MEMLOCK\_FUTURE  
25554 argument to *memlockall()* can be relied upon to cause new shared memory regions to be  
25555 automatically locked.

25556 Existing implementations of *mmap()* return the value `-1` when unsuccessful. Since the casting of  
25557 this value to type `void *` cannot be guaranteed by the ISO C standard to be distinct from a  
25558 successful value, this volume of IEEE Std 1003.1-200x defines the symbol MAP\_FAILED, which a  
25559 conforming implementation does not return as the result of a successful call.

25560 **FUTURE DIRECTIONS**

25561 None.

25562 **SEE ALSO**25563 *exec*, *fcntl()*, *fork()*, *lockf()*, *msync()*, *munmap()*, *mprotect()*, *posix\_typed\_mem\_open()*, *shmat()*,  
25564 *sysconf()*, the Base Definitions volume of IEEE Std 1003.1-200x, <sys/mman.h>25565 **CHANGE HISTORY**

25566 First released in Issue 4, Version 2.

25567 **Issue 5**

25568 Moved from X/OPEN UNIX extension to BASE.

25569 Aligned with *mmap()* in the POSIX Realtime Extension as follows:

- 25570 • The DESCRIPTION is extensively reworded.
- 25571 • The [EAGAIN] and [ENOTSUP] mandatory error conditions are added.
- 25572 • New cases of [ENOMEM] and [ENXIO] are added as mandatory error conditions.
- 25573 • The value returned on failure is the value of the constant MAP\_FAILED; this was previously
- 25574 defined as -1.

25575 Large File Summit extensions are added.

25576 **Issue 6**25577 The *mmap()* function is marked as part of the Memory Mapped Files option.

25578 The Open Group Corrigendum U028/6 is applied, changing (void \*)-1 to MAP\_FAILED.

25579 The following new requirements on POSIX implementations derive from alignment with the  
25580 Single UNIX Specification:

- 25581 • The DESCRIPTION is updated to described the use of MAP\_FIXED.
- 25582 • The DESCRIPTION is updated to describe the addition of an extra reference to the file
- 25583 associated with the file descriptor passed to *mmap()*.
- 25584 • The DESCRIPTION is updated to state that there may be implementation-defined limits on
- 25585 the number of memory regions that can be mapped.
- 25586 • The DESCRIPTION is updated to describe constraints on the alignment and size of the *off*
- 25587 argument.
- 25588 • The [EINVAL] and [EMFILE] error conditions are added.
- 25589 • The [EOVERFLOW] error condition is added. This change is to support large files.

25590 The following changes are made for alignment with the ISO POSIX-1: 1996 standard:

- 25591 • The DESCRIPTION is updated to describe the cases when MAP\_PRIVATE and MAP\_FIXED
- 25592 need not be supported.

25593 The following changes are made for alignment with IEEE Std 1003.1j-2000:

- 25594 • Semantics for typed memory objects are added to the DESCRIPTION.
- 25595 • New [ENOMEM] and [ENXIO] errors are added to the ERRORS section.
- 25596 • The *posix\_typed\_mem\_open()* function is added to the SEE ALSO section.

25597 The DESCRIPTION is updated to avoid use of the term “must” for application requirements.

25598 **NAME**

25599 modf, modff, modfl — decompose a floating-point number

25600 **SYNOPSIS**

25601 #include <math.h>

25602 double modf(double *x*, double \**iptr*);

25603 float modff(float *value*, float \**iptr*);

25604 long double modfl(long double *value*, long double \**iptr*);

25605 **DESCRIPTION**

25606 **CX** The functionality described on this reference page is aligned with the ISO C standard. Any  
25607 conflict between the requirements described here and the ISO C standard is unintentional. This  
25608 volume of IEEE Std 1003.1-200x defers to the ISO C standard.

25609 These functions shall break the argument *x* into integral and fractional parts, each of which has  
25610 the same sign as the argument. It stores the integral part as a **double** (for the *modf()* function), a  
25611 **float** (for the *modff()* function), or a **long double** (for the *modfl()* function), in the object pointed  
25612 to by *iptr*.

25613 **RETURN VALUE**

25614 Upon successful completion, these functions shall return the signed fractional part of *x*.

25615 **MX** If *x* is NaN, a NaN shall be returned, and \**iptr* shall be set to a NaN.

25616 If *x* is  $\pm\text{Inf}$ ,  $\pm 0$  shall be returned, and \**iptr* shall be set to  $\pm\text{Inf}$ .

25617 **ERRORS**

25618 No errors are defined.

25619 **EXAMPLES**

25620 None.

25621 **APPLICATION USAGE**

25622 The *modf()* function computes the function result and \**iptr* such that:

25623 a = modf(x, &iptr) ;

25624 x == a+\*iptr ;

25625 allowing for the usual floating-point inaccuracies.

25626 **RATIONALE**

25627 None.

25628 **FUTURE DIRECTIONS**

25629 None.

25630 **SEE ALSO**

25631 *frexp()*, *isnan()*, *ldexp()*, the Base Definitions volume of IEEE Std 1003.1-200x, <math.h>

25632 **CHANGE HISTORY**

25633 First released in Issue 1. Derived from Issue 1 of the SVID.

25634 **Issue 5**

25635 The DESCRIPTION is updated to indicate how an application should check for an error. This  
25636 text was previously published in the APPLICATION USAGE section.

25637 **Issue 6**

25638 The *modff()* and *modfl()* functions are added for alignment with the ISO/IEC 9899:1999  
25639 standard.



25640 The DESCRIPTION, RETURN VALUE, ERRORS, and APPLICATION USAGE sections are  
25641 revised to align with the ISO/IEC 9899:1999 standard.

25642 IEC 60559:1989 standard floating-point extensions over the ISO/IEC 9899:1999 standard are  
25643 marked.

25644 **NAME**

25645 mprotect — set protection of memory mapping

25646 **SYNOPSIS**

25647 MPR #include &lt;sys/mman.h&gt;

25648 int mprotect(void \*addr, size\_t len, int prot);

25649

25650 **DESCRIPTION**

25651 The *mprotect()* function shall change the access protections to be that specified by *prot* for those  
 25652 whole pages containing any part of the address space of the process starting at address *addr* and  
 25653 continuing for *len* bytes. The parameter *prot* determines whether read, write, execute, or some  
 25654 combination of accesses are permitted to the data being mapped. The *prot* argument should be  
 25655 either PROT\_NONE or the bitwise-inclusive OR of one or more of PROT\_READ, PROT\_WRITE,  
 25656 and PROT\_EXEC.

25657 If an implementation cannot support the combination of access types specified by *prot*, the call  
 25658 to *mprotect()* shall fail.

25659 An implementation may permit accesses other than those specified by *prot*; however, no  
 25660 implementation shall permit a write to succeed where PROT\_WRITE has not been set or shall  
 25661 permit any access where PROT\_NONE alone has been set. Implementations shall support at  
 25662 least the following values of *prot*: PROT\_NONE, PROT\_READ, PROT\_WRITE, and the bitwise-  
 25663 inclusive OR of PROT\_READ and PROT\_WRITE. If PROT\_WRITE is specified, the application  
 25664 shall ensure that it has opened the mapped objects in the specified address range with write  
 25665 permission, unless MAP\_PRIVATE was specified in the original mapping, regardless of whether  
 25666 the file descriptors used to map the objects have since been closed.

25667 The implementation shall require that *addr* be a multiple of the page size as returned by  
 25668 *sysconf()*.

25669 The behavior of this function is unspecified if the mapping was not established by a call to  
 25670 *mmap()*.

25671 When *mprotect()* fails for reasons other than [EINVAL], the protections on some of the pages in  
 25672 the range [*addr,addr+len*) may have been changed.

25673 **RETURN VALUE**

25674 Upon successful completion, *mprotect()* shall return 0; otherwise, it shall return -1 and set *errno*  
 25675 to indicate the error.

25676 **ERRORS**25677 The *mprotect()* function shall fail if:

25678 [EACCES] The *prot* argument specifies a protection that violates the access permission  
 25679 the process has to the underlying memory object.

25680 [EAGAIN] The *prot* argument specifies PROT\_WRITE over a MAP\_PRIVATE mapping  
 25681 and there are insufficient memory resources to reserve for locking the private  
 25682 page.

25683 [EINVAL] The *addr* argument is not a multiple of the page size as returned by *sysconf()*.

25684 [ENOMEM] Addresses in the range [*addr,addr+len*) are invalid for the address space of a  
 25685 process, or specify one or more pages which are not mapped.

25686 [ENOMEM] The *prot* argument specifies PROT\_WRITE on a MAP\_PRIVATE mapping, and  
 25687 it would require more space than the system is able to supply for locking the  
 25688 private pages, if required.

25689 [ENOTSUP] The implementation does not support the combination of accesses requested  
25690 in the *prot* argument.

25691 **EXAMPLES**

25692 None.

25693 **APPLICATION USAGE**

25694 The [EINVAL] error above is marked EX because it is defined as an optional error in the POSIX  
25695 Realtime Extension.

25696 **RATIONALE**

25697 None.

25698 **FUTURE DIRECTIONS**

25699 None.

25700 **SEE ALSO**

25701 *mmap()*, *sysconf()*, the Base Definitions volume of IEEE Std 1003.1-200x, <sys/mman.h>

25702 **CHANGE HISTORY**

25703 First released in Issue 4, Version 2.

25704 **Issue 5**

25705 Moved from X/OPEN UNIX extension to BASE.

25706 Aligned with *mprotect()* in the POSIX Realtime Extension as follows:

- 25707 • The DESCRIPTION is largely reworded.
- 25708 • [ENOTSUP] and a second form of [ENOMEM] are added as mandatory error conditions.
- 25709 • [EAGAIN] is moved from the optional to the mandatory error conditions.

25710 **Issue 6**

25711 The *mprotect()* function is marked as part of the Memory Protection option.

25712 The following new requirements on POSIX implementations derive from alignment with the  
25713 Single UNIX Specification:

- 25714 • The DESCRIPTION is updated to state that implementations require *addr* to be a multiple of  
25715 the page size as returned by *sysconf()*.
- 25716 • The [EINVAL] error condition is added.

25717 The DESCRIPTION is updated to avoid use of the term “must” for application requirements.

25718 **NAME**25719 mq\_close — close a message queue (**REALTIME**)25720 **SYNOPSIS**

25721 MSG #include &lt;mqueue.h&gt;

25722 int mq\_close(mqd\_t mqdes);

25723

25724 **DESCRIPTION**

25725 The *mq\_close()* function shall remove the association between the message queue descriptor,  
25726 *mqdes*, and its message queue. The results of using this message queue descriptor after  
25727 successful return from this *mq\_close()*, and until the return of this message queue descriptor  
25728 from a subsequent *mq\_open()*, are undefined.

25729 If the process has successfully attached a notification request to the message queue via this  
25730 *mqdes*, this attachment shall be removed, and the message queue is available for another process  
25731 to attach for notification.

25732 **RETURN VALUE**

25733 Upon successful completion, the *mq\_close()* function shall return a value of zero; otherwise, the  
25734 function shall return a value of -1 and set *errno* to indicate the error.

25735 **ERRORS**25736 The *mq\_close()* function shall fail if:25737 [EBADF] The *mqdes* argument is not a valid message queue descriptor.25738 **EXAMPLES**

25739 None.

25740 **APPLICATION USAGE**

25741 None.

25742 **RATIONALE**

25743 None.

25744 **FUTURE DIRECTIONS**

25745 None.

25746 **SEE ALSO**

25747 *mq\_open()*, *mq\_unlink()*, *msgctl()*, *msgget()*, *msgrcv()*, *msgsnd()*, the Base Definitions volume of  
25748 IEEE Std 1003.1-200x, <mqueue.h>

25749 **CHANGE HISTORY**

25750 First released in Issue 5. Included for alignment with the POSIX Realtime Extension.

25751 **Issue 6**25752 The *mq\_close()* function is marked as part of the Message Passing option.

25753 The [ENOSYS] error condition has been removed as stubs need not be provided if an  
25754 implementation does not support the Message Passing option.

25755 **NAME**25756 mq\_getattr — get message queue attributes (**REALTIME**)25757 **SYNOPSIS**

25758 MSG #include &lt;mqqueue.h&gt;

25759 int mq\_getattr(mqd\_t mqdes, struct mq\_attr \*mqstat);

25760

25761 **DESCRIPTION**25762 The *mqdes* argument specifies a message queue descriptor.25763 The *mq\_getattr()* function shall obtain status information and attributes of the message queue |  
25764 and the open message queue description associated with the message queue descriptor. |25765 The results shall be returned in the **mq\_attr** structure referenced by the *mqstat* argument. |25766 Upon return, the following members shall have the values associated with the open message |  
25767 queue description as set when the message queue was opened and as modified by subsequent |  
25768 *mq\_setattr()* calls: *mq\_flags*.25769 The following attributes of the message queue shall be returned as set at message queue |  
25770 creation: *mq\_maxmsg*, *mq\_msgsize*.25771 Upon return, the following members within the **mq\_attr** structure referenced by the *mqstat* |  
25772 argument shall be set to the current state of the message queue: |25773 *mq\_curmsgs* The number of messages currently on the queue.25774 **RETURN VALUE**25775 Upon successful completion, the *mq\_getattr()* function shall return zero. Otherwise, the function |  
25776 shall return  $-1$  and set *errno* to indicate the error.25777 **ERRORS**25778 The *mq\_getattr()* function shall fail if:25779 [EBADF] The *mqdes* argument is not a valid message queue descriptor.25780 **EXAMPLES**

25781 None.

25782 **APPLICATION USAGE**

25783 None.

25784 **RATIONALE**

25785 None.

25786 **FUTURE DIRECTIONS**

25787 None.

25788 **SEE ALSO**25789 *mq\_open()*, *mq\_send()*, *mq\_setattr()*, *mq\_timedsend()*, *msgctl()*, *msgget()*, *msgrcv()*, *msgsnd()*, the |  
25790 Base Definitions volume of IEEE Std 1003.1-200x, <mqqueue.h>25791 **CHANGE HISTORY**

25792 First released in Issue 5. Included for alignment with the POSIX Realtime Extension.

25793 **Issue 6**25794 The *mq\_getattr()* function is marked as part of the Message Passing option.25795 The [ENOSYS] error condition has been removed as stubs need not be provided if an |  
25796 implementation does not support the Message Passing option.

25797  
25798

The *mq\_timedsend()* function is added to the SEE ALSO section for alignment with IEEE Std 1003.1d-1999.

25799 **NAME**25800 mq\_notify — notify process that a message is available (**REALTIME**)25801 **SYNOPSIS**

25802 MSG #include &lt;mqqueue.h&gt;

25803 int mq\_notify(mqd\_t mqdes, const struct sigevent \*notification);

25804

25805 **DESCRIPTION**

25806 If the argument *notification* is not NULL, this function shall register the calling process to be |  
 25807 notified of message arrival at an empty message queue associated with the specified message |  
 25808 queue descriptor, *mqdes*. The notification specified by the *notification* argument shall be sent to |  
 25809 the process when the message queue transitions from empty to non-empty. At any time, only |  
 25810 one process may be registered for notification by a message queue. If the calling process or any |  
 25811 other process has already registered for notification of message arrival at the specified message |  
 25812 queue, subsequent attempts to register for that message queue shall fail. |

25813 If *notification* is NULL and the process is currently registered for notification by the specified |  
 25814 message queue, the existing registration shall be removed. |

25815 When the notification is sent to the registered process, its registration shall be removed. The |  
 25816 message queue shall then be available for registration.

25817 If a process has registered for notification of message arrival at a message queue and some |  
 25818 thread is blocked in *mq\_receive()* waiting to receive a message when a message arrives at the |  
 25819 queue, the arriving message shall satisfy the appropriate *mq\_receive()*. The resulting behavior is |  
 25820 as if the message queue remains empty, and no notification shall be sent. |

25821 **RETURN VALUE**

25822 Upon successful completion, the *mq\_notify()* function shall return a value of zero; otherwise, the |  
 25823 function shall return a value of -1 and set *errno* to indicate the error.

25824 **ERRORS**25825 The *mq\_notify()* function shall fail if:25826 [EBADF] The *mqdes* argument is not a valid message queue descriptor.

25827 [EBUSY] A process is already registered for notification by the message queue.

25828 **EXAMPLES**

25829 None.

25830 **APPLICATION USAGE**

25831 None.

25832 **RATIONALE**

25833 None.

25834 **FUTURE DIRECTIONS**

25835 None.

25836 **SEE ALSO**

25837 *mq\_open()*, *mq\_send()*, *mq\_timedsend()*, *msgctl()*, *msgget()*, *msgrcv()*, *msgsnd()*, the Base |  
 25838 Definitions volume of IEEE Std 1003.1-200x, <mqqueue.h>

25839 **CHANGE HISTORY**

25840 First released in Issue 5. Included for alignment with the POSIX Realtime Extension.

25841 **Issue 6**

25842 The *mq\_notify()* function is marked as part of the Message Passing option.

25843 The [ENOSYS] error condition has been removed as stubs need not be provided if an  
25844 implementation does not support the Message Passing option.

25845 The *mq\_timedsend()* function is added to the SEE ALSO section for alignment with  
25846 IEEE Std 1003.1d-1999.



25847 **NAME**25848 `mq_open` — open a message queue (**REALTIME**)25849 **SYNOPSIS**25850 MSG `#include <mqueue.h>`25851 `mqd_t mq_open(const char *name, int oflag, ...);`

25852

25853 **DESCRIPTION**

25854 The `mq_open()` function shall establish the connection between a process and a message queue  
 25855 with a message queue descriptor. It shall create an open message queue description that refers to  
 25856 the message queue, and a message queue descriptor that refers to that open message queue  
 25857 description. The message queue descriptor is used by other functions to refer to that message  
 25858 queue. The *name* argument points to a string naming a message queue. It is unspecified whether  
 25859 the name appears in the file system and is visible to other functions that take pathnames as |  
 25860 arguments. The *name* argument shall conform to the construction rules for a pathname. If *name* |  
 25861 begins with the slash character, then processes calling `mq_open()` with the same value of *name* |  
 25862 shall refer to the same message queue object, as long as that name has not been removed. If *name* |  
 25863 does not begin with the slash character, the effect is implementation-defined. The interpretation  
 25864 of slash characters other than the leading slash character in *name* is implementation-defined. If  
 25865 the *name* argument is not the name of an existing message queue and creation is not requested,  
 25866 `mq_open()` shall fail and return an error.

25867 A message queue descriptor may be implemented using a file descriptor, in which case  
 25868 applications can open up to at least {OPEN\_MAX} file and message queues.

25869 The *oflag* argument requests the desired receive and/or send access to the message queue. The  
 25870 requested access permission to receive messages or send messages shall be granted if the calling |  
 25871 process would be granted read or write access, respectively, to an equivalently protected file. |

25872 The value of *oflag* is the bitwise-inclusive OR of values from the following list. Applications |  
 25873 shall specify exactly one of the first three values (access modes) below in the value of *oflag*: |

25874 **O\_RDONLY** Open the message queue for receiving messages. The process can use the  
 25875 returned message queue descriptor with `mq_receive()`, but not `mq_send()`. A  
 25876 message queue may be open multiple times in the same or different processes  
 25877 for receiving messages.

25878 **O\_WRONLY** Open the queue for sending messages. The process can use the returned  
 25879 message queue descriptor with `mq_send()` but not `mq_receive()`. A message  
 25880 queue may be open multiple times in the same or different processes for  
 25881 sending messages.

25882 **O\_RDWR** Open the queue for both receiving and sending messages. The process can use  
 25883 any of the functions allowed for **O\_RDONLY** and **O\_WRONLY**. A message  
 25884 queue may be open multiple times in the same or different processes for  
 25885 sending messages.

25886 Any combination of the remaining flags may be specified in the value of *oflag*:

25887 **O\_CREAT** Create a message queue. It requires two additional arguments: *mode*, which |  
 25888 shall be of type **mode\_t**, and *attr*, which shall be a pointer to a **mq\_attr** |  
 25889 structure. If the pathname *name* has already been used to create a message |  
 25890 queue that still exists, then this flag shall have no effect, except as noted under |  
 25891 **O\_EXCL**. Otherwise, a message queue shall be created without any messages |  
 25892 in it. The user ID of the message queue shall be set to the effective user ID of |  
 25893 the process, and the group ID of the message queue shall be set to the effective

25894 group ID of the process. The file permission bits shall be set to the value of |  
 25895 *mode*. When bits in *mode* other than file permission bits are set, the effect is |  
 25896 implementation-defined. If *attr* is NULL, the message queue shall be created |  
 25897 with implementation-defined default message queue attributes. If *attr* is non- |  
 25898 NULL and the calling process has the appropriate privilege on *name*, the |  
 25899 message queue *mq\_maxmsg* and *mq\_msgsize* attributes shall be set to the values |  
 25900 of the corresponding members in the **mq\_attr** structure referred to by *attr*. If |  
 25901 *attr* is non-NULL, but the calling process does not have the appropriate |  
 25902 privilege on *name*, the *mq\_open()* function shall fail and return an error |  
 25903 without creating the message queue.

25904 O\_EXCL If O\_EXCL and O\_CREAT are set, *mq\_open()* shall fail if the message queue |  
 25905 *name* exists. The check for the existence of the message queue and the creation |  
 25906 of the message queue if it does not exist shall be atomic with respect to other |  
 25907 threads executing *mq\_open()* naming the same *name* with O\_EXCL and |  
 25908 O\_CREAT set. If O\_EXCL is set and O\_CREAT is not set, the result is |  
 25909 undefined.

25910 O\_NONBLOCK Determines whether a *mq\_send()* or *mq\_receive()* waits for resources or |  
 25911 messages that are not currently available, or fails with *errno* set to [EAGAIN]; |  
 25912 see *mq\_send()* and *mq\_receive()* for details.

25913 The *mq\_open()* function does not add or remove messages from the queue.

25914 **RETURN VALUE**

25915 Upon successful completion, the function shall return a message queue descriptor; otherwise, |  
 25916 the function shall return (**mqd\_t**)-1 and set *errno* to indicate the error.

25917 **ERRORS**

25918 The *mq\_open()* function shall fail if:

25919 [EACCES] The message queue exists and the permissions specified by *oflag* are denied, or |  
 25920 the message queue does not exist and permission to create the message queue |  
 25921 is denied.

25922 [EEXIST] O\_CREAT and O\_EXCL are set and the named message queue already exists.

25923 [EINTR] The *mq\_open()* function was interrupted by a signal.

25924 [EINVAL] The *mq\_open()* function is not supported for the given name.

25925 [EINVAL] O\_CREAT was specified in *oflag*, the value of *attr* is not NULL, and either |  
 25926 *mq\_maxmsg* or *mq\_msgsize* was less than or equal to zero.

25927 [EMFILE] Too many message queue descriptors or file descriptors are currently in use by |  
 25928 this process.

25929 [ENAMETOOLONG]

25930 The length of the *name* argument exceeds {PATH\_MAX} or a pathname |  
 25931 component is longer than {NAME\_MAX}.

25932 [ENFILE] Too many message queues are currently open in the system.

25933 [ENOENT] O\_CREAT is not set and the named message queue does not exist.

25934 [ENOSPC] There is insufficient space for the creation of the new message queue.

25935 **EXAMPLES**

25936 None.

25937 **APPLICATION USAGE**

25938 None.

25939 **RATIONALE**

25940 None.

25941 **FUTURE DIRECTIONS**

25942 None.

25943 **SEE ALSO**

25944 *mq\_close()*, *mq\_getattr()*, *mq\_receive()*, *mq\_send()*, *mq\_setattr()*, *mq\_timedreceive()*, *mq\_timedsend()*,  
25945 *mq\_unlink()*, *msgctl()*, *msgget()*, *msgrcv()*, *msgsnd()*, the Base Definitions volume of  
25946 IEEE Std 1003.1-200x, <mqqueue.h>

25947 **CHANGE HISTORY**

25948 First released in Issue 5. Included for alignment with the POSIX Realtime Extension.

25949 **Issue 6**25950 The *mq\_open()* function is marked as part of the Message Passing option.

25951 The [ENOSYS] error condition has been removed as stubs need not be provided if an  
25952 implementation does not support the Message Passing option.

25953 The *mq\_timedreceive()* and *mq\_timedsend()* functions are added to the SEE ALSO section for  
25954 alignment with IEEE Std 1003.1d-1999.

25955 The DESCRIPTION of O\_EXCL is updated in response to IEEE PASC Interpretation 1003.1c #48.

## 25956 NAME

25957 mq\_receive, mq\_timedreceive — receive a message from a message queue (**REALTIME**)

## 25958 SYNOPSIS

25959 MSG #include <mqueue.h>

```
25960 ssize_t mq_receive(mqd_t mqdes, char *msg_ptr, size_t msg_len,
25961 unsigned *msg_prio);
```

25962

25963 MSG TMO #include <mqueue.h>

25964 #include <time.h>

```
25965 ssize_t mq_timedreceive(mqd_t mqdes, char *restrict msg_ptr,
25966 size_t msg_len, unsigned *restrict msg_prio,
25967 const struct timespec *restrict abs_timeout);
```

25968

## 25969 DESCRIPTION

25970 The *mq\_receive()* function shall receive the oldest of the highest priority message(s) from the  
25971 message queue specified by *mqdes*. If the size of the buffer in bytes, specified by the *msg\_len*  
25972 argument, is less than the *mq\_msgsize* attribute of the message queue, the function shall fail and  
25973 return an error. Otherwise, the selected message shall be removed from the queue and copied to  
25974 the buffer pointed to by the *msg\_ptr* argument.

25975 If the value of *msg\_len* is greater than {SSIZE\_MAX}, the result is implementation-defined.

25976 If the argument *msg\_prio* is not NULL, the priority of the selected message shall be stored in the  
25977 location referenced by *msg\_prio*.

25978 If the specified message queue is empty and O\_NONBLOCK is not set in the message queue  
25979 description associated with *mqdes*, *mq\_receive()* shall block until a message is enqueued on the  
25980 message queue or until *mq\_receive()* is interrupted by a signal. If more than one thread is waiting  
25981 to receive a message when a message arrives at an empty queue and the Priority Scheduling  
25982 option is supported, then the thread of highest priority that has been waiting the longest shall be  
25983 selected to receive the message. Otherwise, it is unspecified which waiting thread receives the  
25984 message. If the specified message queue is empty and O\_NONBLOCK is set in the message  
25985 queue description associated with *mqdes*, no message shall be removed from the queue, and  
25986 *mq\_receive()* shall return an error.

25987 TMO The *mq\_timedreceive()* function shall receive the oldest of the highest priority messages from the  
25988 message queue specified by *mqdes* as described for the *mq\_receive()* function. However, if  
25989 O\_NONBLOCK was not specified when the message queue was opened via the *mq\_open()*  
25990 function, and no message exists on the queue to satisfy the receive, the wait for such a message  
25991 shall be terminated when the specified timeout expires. If O\_NONBLOCK is set, this function is  
25992 equivalent to *mq\_receive()*.

25993 The timeout expires when the absolute time specified by *abs\_timeout* passes, as measured by the  
25994 clock on which timeouts are based (that is, when the value of that clock equals or exceeds  
25995 *abs\_timeout*), or if the absolute time specified by *abs\_timeout* has already been passed at the time  
25996 of the call.

25997 TMO TMR If the Timers option is supported, the timeout shall be based on the CLOCK\_REALTIME clock; if  
25998 the Timers option is not supported, the timeout shall be based on the system clock as returned  
25999 by the *time()* function.

26000 TMO The resolution of the timeout shall be the resolution of the clock on which it is based. The  
26001 *timespec* argument is defined in the <time.h> header.

26002 Under no circumstance shall the operation fail with a timeout if a message can be removed from  
 26003 the message queue immediately. The validity of the *abs\_timeout* parameter need not be checked  
 26004 if a message can be removed from the message queue immediately.

#### 26005 RETURN VALUE

26006 TMO Upon successful completion, the *mq\_receive()* and *mq\_timedreceive()* functions shall return the  
 26007 length of the selected message in bytes and the message shall be removed from the queue.  
 26008 Otherwise, no message shall be removed from the queue, the functions shall return a value of -1,  
 26009 and set *errno* to indicate the error.

#### 26010 ERRORS

26011 TMO The *mq\_receive()* and *mq\_timedreceive()* functions shall fail if:

26012 [EAGAIN] O\_NONBLOCK was set in the message description associated with *mqdes*,  
 26013 and the specified message queue is empty.

26014 [EBADF] The *mqdes* argument is not a valid message queue descriptor open for reading.

26015 [EMSGSIZE] The specified message buffer size, *msg\_len*, is less than the message size  
 26016 attribute of the message queue.

26017 TMO [EINTR] The *mq\_receive()* or *mq\_timedreceive()* operation was interrupted by a signal.

26018 TMO [EINVAL] The process or thread would have blocked, and the *abs\_timeout* parameter  
 26019 specified a nanoseconds field value less than zero or greater than or equal to  
 26020 1 000 million.

26021 TMO [ETIMEDOUT] The O\_NONBLOCK flag was not set when the message queue was opened,  
 26022 but no message arrived on the queue before the specified timeout expired.

26023 TMO The *mq\_receive()* and *mq\_timedreceive()* functions may fail if:

26024 [EBADMSG] The implementation has detected a data corruption problem with the  
 26025 message.

#### 26026 EXAMPLES

26027 None.

#### 26028 APPLICATION USAGE

26029 None.

#### 26030 RATIONALE

26031 None.

#### 26032 FUTURE DIRECTIONS

26033 None.

#### 26034 SEE ALSO

26035 *mq\_open()*, *mq\_send()*, *mq\_timedsend()*, *msgctl()*, *msgget()*, *msgrcv()*, *msgsnd()*, *time()*, the Base  
 26036 Definitions volume of IEEE Std 1003.1-200x, <mqqueue.h>, <time.h>

#### 26037 CHANGE HISTORY

26038 First released in Issue 5. Included for alignment with the POSIX Realtime Extension.

#### 26039 Issue 6

26040 The *mq\_receive()* function is marked as part of the Message Passing option.

26041 The Open Group Corrigendum U021/4 is applied. The DESCRIPTION is changed to refer to  
 26042 *msg\_len* rather than *maxsize*.

26043 The [ENOSYS] error condition has been removed as stubs need not be provided if an  
 26044 implementation does not support the Message Passing option.

26045 The following new requirements on POSIX implementations derive from alignment with the  
26046 Single UNIX Specification:

- 26047 • In this function it is possible for the return value to exceed the range of the type **ssize\_t** (since  
26048 **size\_t** has a larger range of positive values than **ssize\_t**). A sentence restricting the size of  
26049 the **size\_t** object is added to the description to resolve this conflict.

26050 The *mq\_timedreceive()* function is added for alignment with IEEE Std 1003.1d-1999.

26051 The **restrict** keyword is added to the *mq\_timedreceive()* prototype for alignment with the  
26052 ISO/IEC 9899:1999 standard.

26053 IEEE PASC Interpretation 1003.1 #109 is applied, correcting the return type for *mq\_timedreceive()*  
26054 from **int** to **ssize\_t**.

## 26055 NAME

26056 mq\_send, mq\_timedsend — send a message to a message queue (**REALTIME**)

## 26057 SYNOPSIS

26058 MSG #include &lt;mqueue.h&gt;

26059 int mq\_send(mqd\_t mqdes, const char \*msg\_ptr, size\_t msg\_len,  
26060 unsigned msg\_prio);

26061

26062 MSG TMO #include &lt;mqueue.h&gt;

26063 #include &lt;time.h&gt;

26064 int mq\_timedsend(mqd\_t mqdes, const char \*msg\_ptr, size\_t msg\_len,  
26065 unsigned msg\_prio, const struct timespec \*abs\_timeout);

26066

## 26067 DESCRIPTION

26068 The *mq\_send()* function shall add the message pointed to by the argument *msg\_ptr* to the  
26069 message queue specified by *mqdes*. The *msg\_len* argument specifies the length of the message, in  
26070 bytes, pointed to by *msg\_ptr*. The value of *msg\_len* shall be less than or equal to the *mq\_msgsize*  
26071 attribute of the message queue, or *mq\_send()* shall fail.

26072 If the specified message queue is not full, *mq\_send()* shall behave as if the message is inserted  
26073 into the message queue at the position indicated by the *msg\_prio* argument. A message with a  
26074 larger numeric value of *msg\_prio* shall be inserted before messages with lower values of  
26075 *msg\_prio*. A message shall be inserted after other messages in the queue, if any, with equal  
26076 *msg\_prio*. The value of *msg\_prio* shall be less than {MQ\_PRIO\_MAX}.

26077 If the specified message queue is full and O\_NONBLOCK is not set in the message queue  
26078 description associated with *mqdes*, *mq\_send()* shall block until space becomes available to  
26079 enqueue the message, or until *mq\_send()* is interrupted by a signal. If more than one thread is  
26080 waiting to send when space becomes available in the message queue and the Priority Scheduling  
26081 option is supported, then the thread of the highest priority that has been waiting the longest  
26082 shall be unblocked to send its message. Otherwise, it is unspecified which waiting thread is  
26083 unblocked. If the specified message queue is full and O\_NONBLOCK is set in the message  
26084 queue description associated with *mqdes*, the message shall not be queued and *mq\_send()* shall  
26085 return an error.

26086 TMO The *mq\_timedsend()* function shall add a message to the message queue specified by *mqdes* in the  
26087 manner defined for the *mq\_send()* function. However, if the specified message queue is full and  
26088 O\_NONBLOCK is not set in the message queue description associated with *mqdes*, the wait for  
26089 sufficient room in the queue shall be terminated when the specified timeout expires. If  
26090 O\_NONBLOCK is set in the message queue description, this function shall be equivalent to  
26091 *mq\_send()*.

26092 The timeout shall expire when the absolute time specified by *abs\_timeout* passes, as measured by  
26093 the clock on which timeouts are based (that is, when the value of that clock equals or exceeds  
26094 *abs\_timeout*), or if the absolute time specified by *abs\_timeout* has already been passed at the time  
26095 of the call.

26096 TMO TMR If the Timers option is supported, the timeout shall be based on the CLOCK\_REALTIME clock; if  
26097 the Timers option is not supported, the timeout shall be based on the system clock as returned  
26098 by the *time()* function.

26099 TMO The resolution of the timeout shall be the resolution of the clock on which it is based. The  
26100 *timespec* argument is defined in the <time.h> header.

26101 Under no circumstance shall the operation fail with a timeout if there is sufficient room in the  
 26102 queue to add the message immediately. The validity of the *abs\_timeout* parameter need not be  
 26103 checked when there is sufficient room in the queue.

#### 26104 RETURN VALUE

26105 TMO Upon successful completion, the *mq\_send()* and *mq\_timedsend()* functions shall return a value of  
 26106 zero. Otherwise, no message shall be enqueued, the functions shall return  $-1$ , and *errno* shall be  
 26107 set to indicate the error.

#### 26108 ERRORS

26109 TMO The *mq\_send()* and *mq\_timedsend()* functions shall fail if:

26110 [EAGAIN] The O\_NONBLOCK flag is set in the message queue description associated  
 26111 with *mqdes*, and the specified message queue is full.

26112 [EBADF] The *mqdes* argument is not a valid message queue descriptor open for writing.

26113 TMO [EINTR] A signal interrupted the call to *mq\_send()* or *mq\_timedsend()*.

26114 [EINVAL] The value of *msg\_prio* was outside the valid range.

26115 TMO [EINVAL] The process or thread would have blocked, and the *abs\_timeout* parameter  
 26116 specified a nanoseconds field value less than zero or greater than or equal to  
 26117 1 000 million.

26118 [EMSGSIZE] The specified message length, *msg\_len*, exceeds the message size attribute of  
 26119 the message queue.

26120 TMO [ETIMEDOUT] The O\_NONBLOCK flag was not set when the message queue was opened,  
 26121 but the timeout expired before the message could be added to the queue.

#### 26122 EXAMPLES

26123 None.

#### 26124 APPLICATION USAGE

26125 The value of the symbol {MQ\_PRIO\_MAX} limits the number of priority levels supported by the  
 26126 application. Message priorities range from 0 to {MQ\_PRIO\_MAX}-1.

#### 26127 RATIONALE

26128 None.

#### 26129 FUTURE DIRECTIONS

26130 None.

#### 26131 SEE ALSO

26132 *mq\_open()*, *mq\_receive()*, *mq\_setattr()*, *mq\_timedreceive()*, *time()*, the Base Definitions volume of  
 26133 IEEE Std 1003.1-200x, <mqqueue.h>, <time.h>

#### 26134 CHANGE HISTORY

26135 First released in Issue 5. Included for alignment with the POSIX Realtime Extension.

#### 26136 Issue 6

26137 The *mq\_send()* function is marked as part of the Message Passing option.

26138 The [ENOSYS] error condition has been removed as stubs need not be provided if an  
 26139 implementation does not support the Message Passing option.

26140 The *mq\_timedsend()* function is added for alignment with IEEE Std 1003.1d-1999.



26141 **NAME**26142 mq\_setattr — set message queue attributes (**REALTIME**)26143 **SYNOPSIS**

26144 MSG #include &lt;mqqueue.h&gt;

```
26145 int mq_setattr(mqd_t mqdes, const struct mq_attr *restrict mqstat,
26146               struct mq_attr *restrict omqstat);
26147
```

26148 **DESCRIPTION**

26149 The *mq\_setattr()* function shall set attributes associated with the open message queue |  
 26150 description referenced by the message queue descriptor specified by *mqdes*. |

26151 The message queue attributes corresponding to the following members defined in the **mq\_attr** |  
 26152 structure shall be set to the specified values upon successful completion of *mq\_setattr()*: |

26153 *mq\_flags* The value of this member is the bitwise-logical OR of zero or more of |  
 26154 O\_NONBLOCK and any implementation-defined flags.

26155 The values of the *mq\_maxmsg*, *mq\_msgsize*, and *mq\_curmsgs* members of the **mq\_attr** structure |  
 26156 shall be ignored by *mq\_setattr()*.

26157 If *omqstat* is non-NULL, the *mq\_setattr()* function shall store, in the location referenced by |  
 26158 *omqstat*, the previous message queue attributes and the current queue status. These values shall |  
 26159 be the same as would be returned by a call to *mq\_getattr()* at that point. |

26160 **RETURN VALUE**

26161 Upon successful completion, the function shall return a value of zero and the attributes of the |  
 26162 message queue shall have been changed as specified.

26163 Otherwise, the message queue attributes shall be unchanged, and the function shall return a |  
 26164 value of  $-1$  and set *errno* to indicate the error.

26165 **ERRORS**26166 The *mq\_setattr()* function shall fail if:

26167 [EBADF] The *mqdes* argument is not a valid message queue descriptor.

26168 **EXAMPLES**

26169 None.

26170 **APPLICATION USAGE**

26171 None.

26172 **RATIONALE**

26173 None.

26174 **FUTURE DIRECTIONS**

26175 None.

26176 **SEE ALSO**

26177 *mq\_open()*, *mq\_send()*, *mq\_timedsend()*, *msgctl()*, *msgget()*, *msgrcv()*, *msgsnd()*, the Base |  
 26178 Definitions volume of IEEE Std 1003.1-200x, <**mqqueue.h**>

26179 **CHANGE HISTORY**

26180 First released in Issue 5. Included for alignment with the POSIX Realtime Extension.

26181 **Issue 6**

26182 The *mq\_setattr()* function is marked as part of the Message Passing option.

26183 The [ENOSYS] error condition has been removed as stubs need not be provided if an  
26184 implementation does not support the Message Passing option.

26185 The *mq\_timedsend()* function is added to the SEE ALSO section for alignment with  
26186 IEEE Std 1003.1d-1999.

26187 The **restrict** keyword is added to the *mq\_setattr()* prototype for alignment with the  
26188 ISO/IEC 9899:1999 standard.

26189 **NAME**26190 mq\_timedreceive — receive a message from a message queue (**ADVANCED REALTIME**)26191 **SYNOPSIS**

26192 MSG TMO #include &lt;mqueue.h&gt;

26193 #include &lt;time.h&gt;

26194 ssize\_t mq\_timedreceive(mqd\_t mqdes, char \*restrict msg\_ptr,

26195 size\_t msg\_len, unsigned \*restrict msg\_prio,

26196 const struct timespec \*restrict abs\_timeout);

26197

26198 **DESCRIPTION**26199 Refer to *mq\_receive()*.

26200 **NAME**

26201 mq\_timedsend — send a message to a message queue (**ADVANCED REALTIME**)

26202 **SYNOPSIS**

26203 MSG TMO #include <mqueue.h>

26204 #include <time.h>

```
26205 int mq_timedsend(mqd_t mqdes, const char *msg_ptr, size_t msg_len,  
26206                 unsigned msg_prio, const struct timespec *abs_timeout);
```

26207

26208 **DESCRIPTION**

26209 Refer to *mq\_send()*.

26210 **NAME**26211 mq\_unlink — remove a message queue (**REALTIME**)26212 **SYNOPSIS**

26213 MSG #include &lt;mqueue.h&gt;

26214 int mq\_unlink(const char \*name);

26215

26216 **DESCRIPTION**

26217 The *mq\_unlink()* function shall remove the message queue named by the pathname *name*. After |  
 26218 a successful call to *mq\_unlink()* with *name*, a call to *mq\_open()* with *name* shall fail if the flag |  
 26219 O\_CREAT is not set in *flags*. If one or more processes have the message queue open when |  
 26220 *mq\_unlink()* is called, destruction of the message queue shall be postponed until all references to |  
 26221 the message queue have been closed.

26222 Calls to *mq\_open()* to recreate the message queue may fail until the message queue is actually |  
 26223 removed. However, the *mq\_unlink()* call need not block until all references have been closed; it |  
 26224 may return immediately.

26225 **RETURN VALUE**

26226 Upon successful completion, the function shall return a value of zero. Otherwise, the named |  
 26227 message queue shall be unchanged by this function call, and the function shall return a value of |  
 26228 -1 and set *errno* to indicate the error.

26229 **ERRORS**26230 The *mq\_unlink()* function shall fail if:

26231 [EACCES] Permission is denied to unlink the named message queue.

26232 [ENAMETOOLONG]

26233 The length of the *name* argument exceeds {PATH\_MAX} or a pathname |  
 26234 component is longer than {NAME\_MAX}.

26235 [ENOENT] The named message queue does not exist.

26236 **EXAMPLES**

26237 None.

26238 **APPLICATION USAGE**

26239 None.

26240 **RATIONALE**

26241 None.

26242 **FUTURE DIRECTIONS**

26243 None.

26244 **SEE ALSO**

26245 *mq\_close()*, *mq\_open()*, *msgctl()*, *msgget()*, *msgrcv()*, *msgsnd()*, the Base Definitions volume of |  
 26246 IEEE Std 1003.1-200x, <mqueue.h>

26247 **CHANGE HISTORY**

26248 First released in Issue 5. Included for alignment with the POSIX Realtime Extension.

26249 **Issue 6**26250 The *mq\_unlink()* function is marked as part of the Message Passing option.

26251 The Open Group Corrigendum U021/5 is applied, clarifying that upon unsuccessful completion, |  
 26252 the named message queue is unchanged by this function.

26253  
26254

The [ENOSYS] error condition has been removed as stubs need not be provided if an implementation does not support the Message Passing option.

26255 **NAME**

26256 mrand48 — generate uniformly distributed pseudo-random signed long integers

26257 **SYNOPSIS**

26258 xSI #include <stdlib.h>

26259 long mrand48(void);

26260

26261 **DESCRIPTION**

26262 Refer to *drand48()*.

## 26263 NAME

26264 msgctl — XSI message control operations

## 26265 SYNOPSIS

26266 XSI #include &lt;sys/msg.h&gt;

26267 int msgctl(int *msqid*, int *cmd*, struct *msqid\_ds* \**buf*);

26268

## 26269 DESCRIPTION

26270 The *msgctl()* function operates on XSI message queues (see the Base Definitions volume of  
 26271 IEEE Std 1003.1-200x, Section 3.224, Message Queue). It is unspecified whether this function  
 26272 interoperates with the realtime interprocess communication facilities defined in Section 2.8 (on  
 26273 page 491).

26274 The *msgctl()* function shall provide message control operations as specified by *cmd*. The  
 26275 following values for *cmd*, and the message control operations they specify, are:

26276 IPC\_STAT Place the current value of each member of the **msqid\_ds** data structure  
 26277 associated with *msqid* into the structure pointed to by *buf*. The contents of this  
 26278 structure are defined in <sys/msg.h>.

26279 IPC\_SET Set the value of the following members of the **msqid\_ds** data structure  
 26280 associated with *msqid* to the corresponding value found in the structure  
 26281 pointed to by *buf*:

26282 msg\_perm.uid  
 26283 msg\_perm.gid  
 26284 msg\_perm.mode  
 26285 msg\_qbytes

26286 IPC\_SET can only be executed by a process with appropriate privileges or that  
 26287 has an effective user ID equal to the value of **msg\_perm.cuid** or  
 26288 **msg\_perm.uid** in the **msqid\_ds** data structure associated with *msqid*. Only a  
 26289 process with appropriate privileges can raise the value of **msg\_qbytes**.

26290 IPC\_RMID Remove the message queue identifier specified by *msqid* from the system and  
 26291 destroy the message queue and **msqid\_ds** data structure associated with it.  
 26292 IPC\_RMID can only be executed by a process with appropriate privileges or  
 26293 one that has an effective user ID equal to the value of **msg\_perm.cuid** or  
 26294 **msg\_perm.uid** in the **msqid\_ds** data structure associated with *msqid*.

## 26295 RETURN VALUE

26296 Upon successful completion, *msgctl()* shall return 0; otherwise, it shall return -1 and set *errno* to  
 26297 indicate the error.

## 26298 ERRORS

26299 The *msgctl()* function shall fail if:

26300 [EACCES] The argument *cmd* is IPC\_STAT and the calling process does not have read  
 26301 permission; see Section 2.7 (on page 489).

26302 [EINVAL] The value of *msqid* is not a valid message queue identifier; or the value of *cmd*  
 26303 is not a valid command.

26304 [EPERM] The argument *cmd* is IPC\_RMID or IPC\_SET and the effective user ID of the  
 26305 calling process is not equal to that of a process with appropriate privileges  
 26306 and it is not equal to the value of **msg\_perm.cuid** or **msg\_perm.uid** in the data  
 26307 structure associated with *msqid*.



26308 [EPERM] The argument *cmd* is IPC\_SET, an attempt is being made to increase to the  
26309 value of **msg\_qbytes**, and the effective user ID of the calling process does not  
26310 have appropriate privileges.

26311 **EXAMPLES**

26312 None.

26313 **APPLICATION USAGE**

26314 The POSIX Realtime Extension defines alternative interfaces for interprocess communication  
26315 (IPC). Application developers who need to use IPC should design their applications so that  
26316 modules using the IPC routines described in Section 2.7 (on page 489) can be easily modified to  
26317 use the alternative interfaces.

26318 **RATIONALE**

26319 None.

26320 **FUTURE DIRECTIONS**

26321 None.

26322 **SEE ALSO**

26323 *mq\_close()*, *mq\_getattr()*, *mq\_notify()*, *mq\_open()*, *mq\_receive()*, *mq\_send()*, *mq\_setattr()*,  
26324 *mq\_unlink()*, *msgget()*, *msgrcv()*, *msgsnd()*, the Base Definitions volume of IEEE Std 1003.1-200x,  
26325 <sys/msg.h>, Section 2.7 (on page 489)

26326 **CHANGE HISTORY**

26327 First released in Issue 2. Derived from Issue 2 of the SVID.

26328 **Issue 5**

26329 The note about use of POSIX Realtime Extension IPC routines has been moved from FUTURE  
26330 DIRECTIONS to a new APPLICATION USAGE section.

## 26331 NAME

26332 msgget — get the XSI message queue identifier

## 26333 SYNOPSIS

26334 XSI `#include <sys/msg.h>`26335 `int msgget(key_t key, int msgflg);`

26336

## 26337 DESCRIPTION

26338 The *msgget()* function operates on XSI message queues (see the Base Definitions volume of  
 26339 IEEE Std 1003.1-200x, Section 3.224, Message Queue). It is unspecified whether this function  
 26340 interoperates with the realtime interprocess communication facilities defined in Section 2.8 (on  
 26341 page 491).

26342 The *msgget()* function shall return the message queue identifier associated with the argument  
 26343 *key*.

26344 A message queue identifier, associated message queue, and data structure (see `<sys/msg.h>`),  
 26345 shall be created for the argument *key* if one of the following is true:

- 26346 • The argument *key* is equal to `IPC_PRIVATE`.
- 26347 • The argument *key* does not already have a message queue identifier associated with it, and  
 26348 (`msgflg & IPC_CREAT`) is non-zero.

26349 Upon creation, the data structure associated with the new message queue identifier shall be  
 26350 initialized as follows:

- 26351 • `msg_perm.cuid`, `msg_perm.uid`, `msg_perm.cgid`, and `msg_perm.gid` shall be set equal to the  
 26352 effective user ID and effective group ID, respectively, of the calling process.
- 26353 • The low-order 9 bits of `msg_perm.mode` shall be set equal to the low-order 9 bits of *msgflg*.
- 26354 • `msg_qnum`, `msg_lspid`, `msg_lrpid`, `msg_stime`, and `msg_rtime` shall be set equal to 0.
- 26355 • `msg_ctime` shall be set equal to the current time.
- 26356 • `msg_qbytes` shall be set equal to the system limit.

## 26357 RETURN VALUE

26358 Upon successful completion, *msgget()* shall return a non-negative integer, namely a message  
 26359 queue identifier. Otherwise, it shall return `-1` and set *errno* to indicate the error.

## 26360 ERRORS

26361 The *msgget()* function shall fail if:

- |       |          |  |
|-------|----------|--|
| 26362 | [EACCES] | A message queue identifier exists for the argument <i>key</i> , but operation                        |
| 26363 |          | permission as specified by the low-order 9 bits of <i>msgflg</i> would not be granted;               |
| 26364 |          | see Section 2.7 (on page 489).   |
| 26365 | [EEXIST] | A message queue identifier exists for the argument <i>key</i> but <code>((msgflg &amp;</code>        |
| 26366 |          | <code>IPC_CREAT) &amp;&amp; (msgflg &amp; IPC_EXCL)</code> is non-zero.                              |
| 26367 | [ENOENT] | A message queue identifier does not exist for the argument <i>key</i> and <code>(msgflg &amp;</code> |
| 26368 |          | <code>IPC_CREAT)</code> is 0.  |
| 26369 | [ENOSPC] | A message queue identifier is to be created but the system-imposed limit on                          |
| 26370 |          | the maximum number of allowed message queue identifiers system-wide                                  |
| 26371 |          | would be exceeded.   |

**26372 EXAMPLES**

26373 None.

**26374 APPLICATION USAGE**

26375 The POSIX Realtime Extension defines alternative interfaces for interprocess communication  
26376 (IPC). Application developers who need to use IPC should design their applications so that  
26377 modules using the IPC routines described in Section 2.7 (on page 489) can be easily modified to  
26378 use the alternative interfaces.

**26379 RATIONALE**

26380 None.

**26381 FUTURE DIRECTIONS**

26382 None.

**26383 SEE ALSO**

26384 *mq\_close()*, *mq\_getattr()*, *mq\_notify()*, *mq\_open()*, *mq\_receive()*, *mq\_send()*, *mq\_setattr()*,  
26385 *mq\_unlink()*, *msgctl()*, *msgrcv()*, *msgsnd()*, the Base Definitions volume of IEEE Std 1003.1-200x,  
26386 `<sys/msg.h>`, Section 2.7 (on page 489)

**26387 CHANGE HISTORY**

26388 First released in Issue 2. Derived from Issue 2 of the SVID.

**26389 Issue 5**

26390 The note about use of POSIX Realtime Extension IPC routines has been moved from FUTURE  
26391 DIRECTIONS to a new APPLICATION USAGE section.

## 26392 NAME

26393 msgrcv — XSI message receive operation

## 26394 SYNOPSIS

26395 XSI 

```
#include <sys/msg.h>
```

```
26396 ssize_t msgrcv(int msqid, void *msgp, size_t msgsz, long msgtyp,
26397               int msgflg);
26398
```

## 26399 DESCRIPTION

26400 The *msgrcv()* function operates on XSI message queues (see the Base Definitions volume of  
 26401 IEEE Std 1003.1-200x, Section 3.224, Message Queue). It is unspecified whether this function  
 26402 interoperates with the realtime interprocess communication facilities defined in Section 2.8 (on  
 26403 page 491).

26404 The *msgrcv()* function shall read a message from the queue associated with the message queue  
 26405 identifier specified by *msqid* and place it in the user-defined buffer pointed to by *msgp*.

26406 The application shall ensure that the argument *msgp* points to a user-defined buffer that contains  
 26407 first a field of type **long** specifying the type of the message, and then a data portion that holds  
 26408 the data bytes of the message. The structure below is an example of what this user-defined  
 26409 buffer might look like:

```
26410 struct mymsg {
26411     long   mtype;      /* Message type. */
26412     char   mtext[1];  /* Message text. */
26413 }
```

26414 The structure member *mtype* is the received message's type as specified by the sending process.

26415 The structure member *mtext* is the text of the message.

26416 The argument *msgsz* specifies the size in bytes of *mtext*. The received message shall be truncated  
 26417 to *msgsz* bytes if it is larger than *msgsz* and (*msgflg* & MSG\_NOERROR) is non-zero. The  
 26418 truncated part of the message shall be lost and no indication of the truncation shall be given to  
 26419 the calling process.

26420 If the value of *msgsz* is greater than {SSIZE\_MAX}, the result is implementation-defined.

26421 The argument *msgtyp* specifies the type of message requested as follows:

- 26422 • If *msgtyp* is 0, the first message on the queue shall be received.
- 26423 • If *msgtyp* is greater than 0, the first message of type *msgtyp* shall be received.
- 26424 • If *msgtyp* is less than 0, the first message of the lowest type that is less than or equal to the  
 26425 absolute value of *msgtyp* shall be received.

26426 The argument *msgflg* specifies the action to be taken if a message of the desired type is not on the  
 26427 queue. These are as follows:

- 26428 • If (*msgflg* & IPC\_NOWAIT) is non-zero, the calling thread shall return immediately with a  
 26429 return value of -1 and *errno* set to [ENOMSG].
- 26430 • If (*msgflg* & IPC\_NOWAIT) is 0, the calling thread shall suspend execution until one of the  
 26431 following occurs:
  - 26432 — A message of the desired type is placed on the queue.
  - 26433 — The message queue identifier *msqid* is removed from the system; when this occurs, *errno*  
 26434 shall be set equal to [EIDRM] and -1 shall be returned.

26435 — The calling thread receives a signal that is to be caught; in this case a message is not  
 26436 received and the calling thread resumes execution in the manner prescribed in *sigaction()*.

26437 Upon successful completion, the following actions are taken with respect to the data structure  
 26438 associated with *msqid*:

- 26439 • **msg\_qnum** shall be decremented by 1. |
- 26440 • **msg\_lrpid** shall be set equal to the process ID of the calling process. |
- 26441 • **msg\_rtime** shall be set equal to the current time. |

#### 26442 RETURN VALUE

26443 Upon successful completion, *msgrcv()* shall return a value equal to the number of bytes actually  
 26444 placed into the buffer *mtext*. Otherwise, no message shall be received, *msgrcv()* shall return  
 26445 (**ssize\_t**)−1, and *errno* shall be set to indicate the error.

#### 26446 ERRORS

26447 The *msgrcv()* function shall fail if:

- |                |          |   |
|----------------|----------|---|
| 26448          | [E2BIG]  | The value of <i>mtext</i> is greater than <i>msgsz</i> and ( <i>msgflg</i> & MSG_NOERROR) is 0.         |
| 26449<br>26450 | [EACCES] | Operation permission is denied to the calling process; see Section 2.7 (on page 489).                   |
| 26451          | [EIDRM]  | The message queue identifier <i>msqid</i> is removed from the system.                                   |
| 26452          | [EINTR]  | The <i>msgrcv()</i> function was interrupted by a signal.   |
| 26453          | [EINVAL] | <i>msqid</i> is not a valid message queue identifier.   |
| 26454<br>26455 | [ENOMSG] | The queue does not contain a message of the desired type and ( <i>msgflg</i> & IPC_NOWAIT) is non-zero. |

#### 26456 EXAMPLES

##### 26457 Receiving a Message

26458 The following example receives the first message on the queue (based on the value of the *msgtyp*  
 26459 argument, 0). The queue is identified by the *msqid* argument (assuming that the value has  
 26460 previously been set). This call specifies that an error should be reported if no message is  
 26461 available, but not if the message is too large. The message size is calculated directly using the  
 26462 *sizeof* operator.

```

26463 #include <sys/msg.h>
26464 ...
26465 int result;
26466 int msqid;
26467 struct message {
26468     long type;
26469     char text[20];
26470 } msg;
26471 long msgtyp = 0;
26472 ...
26473 result = msgrcv(msqid, (void *) &msg, sizeof(msg.text),
26474               msgtyp, MSG_NOERROR | IPC_NOWAIT);
  
```

**26475 APPLICATION USAGE**

26476 The POSIX Realtime Extension defines alternative interfaces for interprocess communication  
26477 (IPC). Application developers who need to use IPC should design their applications so that  
26478 modules using the IPC routines described in Section 2.7 (on page 489) can be easily modified to  
26479 use the alternative interfaces.

**26480 RATIONALE**

26481 None.

**26482 FUTURE DIRECTIONS**

26483 None.

**26484 SEE ALSO**

26485 *mq\_close()*, *mq\_getattr()*, *mq\_notify()*, *mq\_open()*, *mq\_receive()*, *mq\_send()*, *mq\_setattr()*,  
26486 *mq\_unlink()*, *msgctl()*, *msgget()*, *msgsnd()*, *sigaction()*, the Base Definitions volume of  
26487 IEEE Std 1003.1-200x, <sys/msg.h>, Section 2.7 (on page 489)

**26488 CHANGE HISTORY**

26489 First released in Issue 2. Derived from Issue 2 of the SVID.

**26490 Issue 5**

26491 The type of the return value is changed from **int** to **ssize\_t**, and a warning is added to the  
26492 DESCRIPTION about values of *msgsz* larger than {SSIZE\_MAX}.

26493 The note about use of POSIX Realtime Extension IPC routines has been moved from FUTURE  
26494 DIRECTIONS to the APPLICATION USAGE section.

**26495 Issue 6**

26496 The DESCRIPTION is updated to avoid use of the term “must” for application requirements.

26497 **NAME**

26498 msgsnd — XSI message send operation

26499 **SYNOPSIS**26500 XSI 

```
#include <sys/msg.h>
```

26501 

```
int msgsnd(int msqid, const void *msgp, size_t msgsz, int msgflg);
```

26502

26503 **DESCRIPTION**

26504 The *msgsnd()* function operates on XSI message queues (see the Base Definitions volume of  
 26505 IEEE Std 1003.1-200x, Section 3.224, Message Queue). It is unspecified whether this function  
 26506 interoperates with the realtime interprocess communication facilities defined in Section 2.8 (on  
 26507 page 491).

26508 The *msgsnd()* function shall send a message to the queue associated with the message queue  
 26509 identifier specified by *msqid*.

26510 The application shall ensure that the argument *msgp* points to a user-defined buffer that contains  
 26511 first a field of type **long** specifying the type of the message, and then a data portion that holds  
 26512 the data bytes of the message. The structure below is an example of what this user-defined  
 26513 buffer might look like:

```
26514 struct mymsg {
26515     long   mtype;           /* Message type. */
26516     char   mtext[1];       /* Message text. */
26517 }
```

26518 The structure member *mtype* is a non-zero positive type **long** that can be used by the receiving  
 26519 process for message selection.

26520 The structure member *mtext* is any text of length *msgsz* bytes. The argument *msgsz* can range  
 26521 from 0 to a system-imposed maximum.

26522 The argument *msgflg* specifies the action to be taken if one or more of the following are true:

- 26523 • The number of bytes already on the queue is equal to **msg\_qbytes**; see `<sys/msg.h>`.
- 26524 • The total number of messages on all queues system-wide is equal to the system-imposed  
 26525 limit.

26526 These actions are as follows:

- 26527 • If (*msgflg* & IPC\_NOWAIT) is non-zero, the message shall not be sent and the calling thread  
 26528 shall return immediately.
- 26529 • If (*msgflg* & IPC\_NOWAIT) is 0, the calling thread shall suspend execution until one of the  
 26530 following occurs:
  - 26531 — The condition responsible for the suspension no longer exists, in which case the message  
 26532 is sent.
  - 26533 — The message queue identifier *msqid* is removed from the system; when this occurs, *errno*  
 26534 shall be set equal to [EIDRM] and `-1` shall be returned.
  - 26535 — The calling thread receives a signal that is to be caught; in this case the message is not  
 26536 sent and the calling thread resumes execution in the manner prescribed in *sigaction()*.

26537 Upon successful completion, the following actions are taken with respect to the data structure  
 26538 associated with *msqid*; see `<sys/msg.h>`:

- 26539           • **msg\_qnum** shall be incremented by 1. |
- 26540           • **msg\_lspid** shall be set equal to the process ID of the calling process. |
- 26541           • **msg\_stime** shall be set equal to the current time. |

**26542 RETURN VALUE**

26543           Upon successful completion, *msgsnd()* shall return 0; otherwise, no message shall be sent,  
26544           *msgsnd()* shall return -1, and *errno* shall be set to indicate the error.

**26545 ERRORS**

26546           The *msgsnd()* function shall fail if:

- 26547           [EACCES]           Operation permission is denied to the calling process; see Section 2.7 (on page  
26548                                   489).
- 26549           [EAGAIN]           The message cannot be sent for one of the reasons cited above and (*msgflg* &  
26550                                   IPC\_NOWAIT) is non-zero.
- 26551           [EIDRM]           The message queue identifier *msqid* is removed from the system.
- 26552           [EINTR]           The *msgsnd()* function was interrupted by a signal.
- 26553           [EINVAL]           The value of *msqid* is not a valid message queue identifier, or the value of  
26554                                   *mtype* is less than 1; or the value of *msgsz* is less than 0 or greater than the  
26555                                   system-imposed limit.

**26556 EXAMPLES****26557 Sending a Message**

26558           The following example sends a message to the queue identified by the *msqid* argument  
26559           (assuming that value has previously been set). This call specifies that an error should be  
26560           reported if no message is available. The message size is calculated directly using the *sizeof*  
26561           operator.

```
26562           #include <sys/msg.h>
26563           ...
26564           int result;
26565           int msqid;
26566           struct message {
26567               long type;
26568               char text[20];
26569           } msg;

26570           msg.type = 1;
26571           strcpy(msg.text, "This is message 1");
26572           ...
26573           result = msgsnd(msqid, (void *) &msg, sizeof(msg.text), IPC_NOWAIT);
```

**26574 APPLICATION USAGE**

26575           The POSIX Realtime Extension defines alternative interfaces for interprocess communication  
26576           (IPC). Application developers who need to use IPC should design their applications so that  
26577           modules using the IPC routines described in Section 2.7 (on page 489) can be easily modified to  
26578           use the alternative interfaces.



**26579 RATIONALE**

26580 None.

**26581 FUTURE DIRECTIONS**

26582 None.

**26583 SEE ALSO**

26584 *mq\_close()*, *mq\_getattr()*, *mq\_notify()*, *mq\_open()*, *mq\_receive()*, *mq\_send()*, *mq\_setattr()*,  
26585 *mq\_unlink()*, *msgctl()*, *msgget()*, *msgrcv()*, *sigaction()*, the Base Definitions volume of  
26586 IEEE Std 1003.1-200x, <sys/msg.h>, Section 2.7 (on page 489)

**26587 CHANGE HISTORY**

26588 First released in Issue 2. Derived from Issue 2 of the SVID.

**26589 Issue 5**

26590 The note about use of POSIX Realtime Extension IPC routines has been moved from FUTURE  
26591 DIRECTIONS to a new APPLICATION USAGE section.

**26592 Issue 6**

26593 The DESCRIPTION is updated to avoid use of the term “must” for application requirements.

26594 **NAME**

26595       msync — synchronize memory with physical storage

26596 **SYNOPSIS**

26597 MF SIO   #include &lt;sys/mman.h&gt;

26598       int msync(void \*addr, size\_t len, int flags);

26599

26600 **DESCRIPTION**

26601       The *msync()* function shall write all modified data to permanent storage locations, if any, in  
 26602       those whole pages containing any part of the address space of the process starting at address  
 26603       *addr* and continuing for *len* bytes. If no such storage exists, *msync()* need not have any effect. If  
 26604       requested, the *msync()* function shall then invalidate cached copies of data.

26605       The implementation shall require that *addr* be a multiple of the page size as returned by  
 26606       *sysconf()*.

26607       For mappings to files, the *msync()* function shall ensure that all write operations are completed  
 26608       as defined for synchronized I/O data integrity completion. It is unspecified whether the  
 26609       implementation also writes out other file attributes. When the *msync()* function is called on  
 26610       MAP\_PRIVATE mappings, any modified data shall not be written to the underlying object and  
 26611       shall not cause such data to be made visible to other processes. It is unspecified whether data in  
 26612 SHM|TYM MAP\_PRIVATE mappings has any permanent storage locations. The effect of *msync()* on a  
 26613       shared memory object or a typed memory object is unspecified. The behavior of this function is  
 26614       unspecified if the mapping was not established by a call to *mmap()*.

26615       The *flags* argument is constructed from the bitwise-inclusive OR of one or more of the following  
 26616       flags defined in the <sys/mman.h> header:

26617

26618

26619

26620

26621

Symbolic Constant	Description
MS_ASYNC	Perform asynchronous writes.
MS_SYNC	Perform synchronous writes.
MS_INVALIDATE	Invalidate cached data.

26622       When MS\_ASYNC is specified, *msync()* shall return immediately once all the write operations  
 26623       are initiated or queued for servicing; when MS\_SYNC is specified, *msync()* shall not return until  
 26624       all write operations are completed as defined for synchronized I/O data integrity completion.  
 26625       Either MS\_ASYNC or MS\_SYNC is specified, but not both.

26626       When MS\_INVALIDATE is specified, *msync()* shall invalidate all cached copies of mapped data  
 26627       that are inconsistent with the permanent storage locations such that subsequent references shall  
 26628       obtain data that was consistent with the permanent storage locations sometime between the call  
 26629       to *msync()* and the first subsequent memory reference to the data.

26630       If *msync()* causes any write to a file, the file's *st\_ctime* and *st\_mtime* fields shall be marked for  
 26631       update.

26632 **RETURN VALUE**

26633       Upon successful completion, *msync()* shall return 0; otherwise, it shall return -1 and set *errno* to  
 26634       indicate the error.

26635 **ERRORS**26636       The *msync()* function shall fail if:

26637       [EBUSY]       Some or all of the addresses in the range starting at *addr* and continuing for *len*  
 26638       bytes are locked, and MS\_INVALIDATE is specified.

- 26639 [EINVAL] The value of *flags* is invalid.
- 26640 [EINVAL] The value of *addr* is not a multiple of the page size, {PAGESIZE}.
- 26641 [ENOMEM] The addresses in the range starting at *addr* and continuing for *len* bytes are outside the range allowed for the address space of a process or specify one or more pages that are not mapped.
- 26642
- 26643

26644 **EXAMPLES**

26645 None.

26646 **APPLICATION USAGE**

26647 The *msync()* function is only supported if the Memory Mapped Files option and the Synchronized Input and Output option are supported, and thus need not be available on all implementations.

26648

26649

26650 The *msync()* function should be used by programs that require a memory object to be in a known state; for example, in building transaction facilities.

26651

26652 Normal system activity can cause pages to be written to disk. Therefore, there are no guarantees that *msync()* is the only control over when pages are or are not written to disk.

26653

26654 **RATIONALE**

26655 The *msync()* function writes out data in a mapped region to the permanent storage for the underlying object. The call to *msync()* ensures data integrity of the file.

26656

26657 After the data is written out, any cached data may be invalidated if the MS\_INVALIDATE flag was specified. This is useful on systems that do not support read/write consistency.

26658

26659 **FUTURE DIRECTIONS**

26660 None.

26661 **SEE ALSO**26662 *mmap()*, *sysconf()*, the Base Definitions volume of IEEE Std 1003.1-200x, <sys/mman.h>26663 **CHANGE HISTORY**

26664 First released in Issue 4, Version 2.

26665 **Issue 5**

26666 Moved from X/OPEN UNIX extension to BASE.

26667 Aligned with *msync()* in the POSIX Realtime Extension as follows:

- 26668
- The DESCRIPTION is extensively reworded.
  - [EBUSY] and a new form of [EINVAL] are added as mandatory error conditions.
- 26669

26670 **Issue 6**

26671 The *msync()* function is marked as part of the Memory Mapped Files and Synchronized Input and Output options.

26672

26673 The following changes are made for alignment with the ISO POSIX-1: 1996 standard:

- 26674
- The [EBUSY] mandatory error condition is added.

26675 The following new requirements on POSIX implementations derive from alignment with the Single UNIX Specification:

26676

- 26677
- The DESCRIPTION is updated to state that implementations require *addr* to be a multiple of the page size.
  - The second [EINVAL] error condition is made mandatory.
- 26678
- 26679

26680  
26681

The DESCRIPTION is updated for alignment with IEEE Std 1003.1j-2000 by adding reference to typed memory objects.

26682 **NAME**

26683       munlock — unlock a range of process address space

26684 **SYNOPSIS**

26685 MLR       #include &lt;sys/mman.h&gt;

26686       int munlock(const void \*addr, size\_t len);

26687

26688 **DESCRIPTION**26689       Refer to *mlock()*.

26690 **NAME**

26691           munlockall — unlock the address space of a process

26692 **SYNOPSIS**

26693 ML       #include <sys/mman.h>

26694           int munlockall(void);

26695

26696 **DESCRIPTION**

26697           Refer to *mlockall()*.

26698 **NAME**26699 `munmap` — unmap pages of memory26700 **SYNOPSIS**26701 MF|SHM `#include <sys/mman.h>`26702 `int munmap(void *addr, size_t len);`

26703

26704 **DESCRIPTION**

26705 The `munmap()` function shall remove any mappings for those entire pages containing any part of  
 26706 the address space of the process starting at `addr` and continuing for `len` bytes. Further references |  
 26707 to these pages shall result in the generation of a SIGSEGV signal to the process. If there are no |  
 26708 mappings in the specified address range, then `munmap()` has no effect.

26709 The implementation shall require that `addr` be a multiple of the page size {PAGESIZE}.

26710 If a mapping to be removed was private, any modifications made in this address range shall be  
 26711 discarded.

26712 ML|MLR Any memory locks (see `mlock()` and `mlockall()`) associated with this address range shall be  
 26713 removed, as if by an appropriate call to `munlock()`.

26714 TYM If a mapping removed from a typed memory object causes the corresponding address range of  
 26715 the memory pool to be inaccessible by any process in the system except through allocatable  
 26716 mappings (that is, mappings of typed memory objects opened with the  
 26717 POSIX\_TYPED\_MEM\_MAP\_ALLOCATABLE flag), then that range of the memory pool shall  
 26718 become deallocated and may become available to satisfy future typed memory allocation  
 26719 requests.

26720 A mapping removed from a typed memory object opened with the  
 26721 POSIX\_TYPED\_MEM\_MAP\_ALLOCATABLE flag shall not affect in any way the availability of  
 26722 that typed memory for allocation.

26723 The behavior of this function is unspecified if the mapping was not established by a call to  
 26724 `mmap()`.

26725 **RETURN VALUE**

26726 Upon successful completion, `munmap()` shall return 0; otherwise, it shall return `-1` and set `errno`  
 26727 to indicate the error.

26728 **ERRORS**

26729 The `munmap()` function shall fail if:

26730 [EINVAL] Addresses in the range `[addr,addr+len)` are outside the valid range for the  
 26731 address space of a process.

26732 [EINVAL] The `len` argument is 0.

26733 [EINVAL] The `addr` argument is not a multiple of the page size as returned by `sysconf()`.

26734 **EXAMPLES**

26735 None.

26736 **APPLICATION USAGE**

26737 The *munmap()* function is only supported if the Memory Mapped Files option or the Shared  
26738 Memory Objects option is supported.

26739 **RATIONALE**26740 The *munmap()* function corresponds to SVR4, just as the *mmap()* function does.

26741 It is possible that an application has applied process memory locking to a region that contains  
26742 shared memory. If this has occurred, the *munmap()* call ignores those locks and, if necessary,  
26743 causes those locks to be removed.

26744 **FUTURE DIRECTIONS**

26745 None.

26746 **SEE ALSO**

26747 *mlock()*, *mlockall()*, *mmap()*, *posix\_typed\_mem\_open()*, *sysconf()*, the Base Definitions volume of  
26748 IEEE Std 1003.1-200x, <signal.h>, <sys/mman.h>

26749 **CHANGE HISTORY**

26750 First released in Issue 4, Version 2.

26751 **Issue 5**

26752 Moved from X/OPEN UNIX extension to BASE.

26753 Aligned with *munmap()* in the POSIX Realtime Extension as follows:

- 26754 • The DESCRIPTION is extensively reworded.
- 26755 • The SIGBUS error is no longer permitted to be generated.

26756 **Issue 6**

26757 The *munmap()* function is marked as part of the Memory Mapped Files and Shared Memory  
26758 Objects option.

26759 The following new requirements on POSIX implementations derive from alignment with the  
26760 Single UNIX Specification:

- 26761 • The DESCRIPTION is updated to state that implementations require *addr* to be a multiple of  
26762 the page size.
- 26763 • The [EINVAL] error conditions are added.

26764 The following changes are made for alignment with IEEE Std 1003.1j-2000:

- 26765 • Semantics for typed memory objects are added to the DESCRIPTION.
- 26766 • The *posix\_typed\_mem\_open()* function is added to the SEE ALSO section.



26767 **NAME**

26768 nan, nanf, nanl — return quiet NaN

26769 **SYNOPSIS**

26770 #include &lt;math.h&gt;

26771 double nan(const char \*tagp);

26772 float nanf(const char \*tagp);

26773 long double nanl(const char \*tagp);

26774 **DESCRIPTION**

26775 cx The functionality described on this reference page is aligned with the ISO C standard. Any  
 26776 conflict between the requirements described here and the ISO C standard is unintentional. This  
 26777 volume of IEEE Std 1003.1-200x defers to the ISO C standard.

26778 The function call *nan*("n-char-sequence") shall be equivalent to:

26779 strtod("NAN(n-char-sequence)", (char \*\*) NULL);

26780 The function call *nan*("") shall be equivalent to:

26781 strtod("NAN()", (char \*\*) NULL)

26782 If *tagp* does not point to an *n-char* sequence or an empty string, the function call shall be  
 26783 equivalent to:

26784 strtod("NAN", (char \*\*) NULL)

26785 Function calls to *nanf*() and *nanl*() are equivalent to the corresponding function calls to *strtof*()  
 26786 and *strtold*().

26787 **RETURN VALUE**26788 These functions shall return a quiet NaN, if available, with content indicated through *tagp*.

26789 If the implementation does not support quiet NaNs, these functions shall return zero.

26790 **ERRORS**

26791 No errors are defined.

26792 **EXAMPLES**

26793 None.

26794 **APPLICATION USAGE**

26795 None.

26796 **RATIONALE**

26797 None.

26798 **FUTURE DIRECTIONS**

26799 None.

26800 **SEE ALSO**26801 *strtod*(), *strtold*(), the Base Definitions volume of IEEE Std 1003.1-200x, <math.h>26802 **CHANGE HISTORY**

26803 First released in Issue 6. Derived from the ISO/IEC 9899:1999 standard.

26804 **NAME**26805 nanosleep — high resolution sleep (**REALTIME**)26806 **SYNOPSIS**26807 TMR `#include <time.h>`26808 `int nanosleep(const struct timespec *rqtp, struct timespec *rmtp);`

26809

26810 **DESCRIPTION**26811 The *nanosleep()* function shall cause the current thread to be suspended from execution until  
26812 either the time interval specified by the *rqtp* argument has elapsed or a signal is delivered to the  
26813 calling thread, and its action is to invoke a signal-catching function or to terminate the process.26814 The suspension time may be longer than requested because the argument value is rounded up to  
26815 an integer multiple of the sleep resolution or because of the scheduling of other activity by the  
26816 system. But, except for the case of being interrupted by a signal, the suspension time shall not be  
26817 less than the time specified by *rqtp*, as measured by the system clock, `CLOCK_REALTIME`.26818 The use of the *nanosleep()* function has no effect on the action or blockage of any signal.26819 **RETURN VALUE**26820 If the *nanosleep()* function returns because the requested time has elapsed, its return value shall  
26821 be zero.26822 If the *nanosleep()* function returns because it has been interrupted by a signal, it shall return a  
26823 value of `-1` and set *errno* to indicate the interruption. If the *rmtp* argument is non-NULL, the  
26824 **timespec** structure referenced by it is updated to contain the amount of time remaining in the  
26825 interval (the requested time minus the time actually slept). If the *rmtp* argument is NULL, the  
26826 remaining time is not returned.26827 If *nanosleep()* fails, it shall return a value of `-1` and set *errno* to indicate the error.26828 **ERRORS**26829 The *nanosleep()* function shall fail if:26830 [EINTR] The *nanosleep()* function was interrupted by a signal.26831 [EINVAL] The *rqtp* argument specified a nanosecond value less than zero or greater than  
26832 or equal to 1000 million.26833 **EXAMPLES**

26834 None.

26835 **APPLICATION USAGE**

26836 None.

26837 **RATIONALE**26838 It is common to suspend execution of a process for an interval in order to poll the status of a  
26839 non-interrupting function. A large number of actual needs can be met with a simple extension to  
26840 *sleep()* that provides finer resolution.26841 In the POSIX.1-1990 standard and SVR4, it is possible to implement such a routine, but the  
26842 frequency of wakeup is limited by the resolution of the *alarm()* and *sleep()* functions. In 4.3 BSD,  
26843 it is possible to write such a routine using no static storage and reserving no system facilities.  
26844 Although it is possible to write a function with similar functionality to *sleep()* using the  
26845 remainder of the timers function, such a function requires the use of signals and the reservation  
26846 of some signal number. This volume of IEEE Std 1003.1-200x requires that *nanosleep()* be non-  
26847 intrusive of the signals function.

26848           The *nanosleep()* function shall return a value of 0 on success and –1 on failure or if interrupted.  
26849           This latter case is different from *sleep()*. This was done because the remaining time is returned  
26850           via an argument structure pointer, *rmtp*, instead of as the return value.

26851 **FUTURE DIRECTIONS**

26852           None.

26853 **SEE ALSO**

26854           *sleep()*, the Base Definitions volume of IEEE Std 1003.1-200x, <**time.h**>

26855 **CHANGE HISTORY**

26856           First released in Issue 5. Included for alignment with the POSIX Realtime Extension.

26857 **Issue 6**

26858           The *nanosleep()* function is marked as part of the Timers option.

26859           The [ENOSYS] error condition has been removed as stubs need not be provided if an  
26860           implementation does not support the Timers option.

26861 **NAME**

26862 nearbyint, nearbyintf, nearbyintl — floating-point rounding functions

26863 **SYNOPSIS**

26864 #include &lt;math.h&gt;

26865 double nearbyint(double x);

26866 float nearbyintf(float x);

26867 long double nearbyintl(long double x);

26868 **DESCRIPTION**

26869 CX The functionality described on this reference page is aligned with the ISO C standard. Any  
 26870 conflict between the requirements described here and the ISO C standard is unintentional. This  
 26871 volume of IEEE Std 1003.1-200x defers to the ISO C standard.

26872 These functions shall round their argument to an integer value in floating-point format, using  
 26873 the current rounding direction and without raising the inexact floating-point exception.

26874 An application wishing to check for error situations should set *errno* to zero and call  
 26875 *feclearexcept*(FE\_ALL\_EXCEPT) before calling these functions. On return, if *errno* is non-zero or  
 26876 *fetestexcept*(FE\_INVALID | FE\_DIVBYZERO | FE\_OVERFLOW | FE\_UNDERFLOW) is non-  
 26877 zero, an error has occurred.

26878 **RETURN VALUE**

26879 Upon successful completion, these functions shall return the rounded integer value.

26880 MX If *x* is NaN, a NaN shall be returned.26881 If *x* is  $\pm 0$ ,  $\pm 0$  shall be returned.26882 If *x* is  $\pm \text{Inf}$ , *x* shall be returned.

26883 XSI If the correct value would cause overflow, a range error shall occur and *nearbyint*(), *nearbyintf*(),  
 26884 and *nearbyintl*() shall return the value of the macro HUGE\_VAL, HUGE\_VALF, and  
 26885 HUGE\_VALL, respectively.

26886 **ERRORS**

26887 These functions shall fail if:

26888 XSI **Range Error** The result would cause an overflow.

26889 If the integer expression (math\_errhandling & MATH\_ERRNO) is non-zero, |  
 26890 then *errno* shall be set to [ERANGE]. If the integer expression |  
 26891 (math\_errhandling & MATH\_ERREXCEPT) is non-zero, then the overflow |  
 26892 floating-point exception shall be raised. |

26893 **EXAMPLES**

26894 None.

26895 **APPLICATION USAGE**

26896 On error, the expressions (math\_errhandling & MATH\_ERRNO) and (math\_errhandling &  
 26897 MATH\_ERREXCEPT) are independent of each other, but at least one of them must be non-zero.

26898 **RATIONALE**

26899 None.

26900 **FUTURE DIRECTIONS**

26901 None.

26902 **SEE ALSO**

26903            *feclearexcept()*, *fetestexcept()*, the Base Definitions volume of IEEE Std 1003.1-200x, Section 4.18, |  
26904            Treatment of Error Conditions for Mathematical Functions, <math.h> |

26905 **CHANGE HISTORY**

26906            First released in Issue 6. Derived from the ISO/IEC 9899:1999 standard.

## 26907 NAME

26908 nextafter, nextafterf, nextafterl, nexttoward, nexttowardf, nexttowardl — next representable  
26909 floating-point number

## 26910 SYNOPSIS

```
26911 #include <math.h>

26912 double nextafter(double x, double y);
26913 float nextafterf(float x, float y);
26914 long double nextafterl(long double x, long double y);
26915 double nexttoward(double x, long double y);
26916 float nexttowardf(float x, long double y);
26917 long double nexttowardl(long double x, long double y);
```

## 26918 DESCRIPTION

26919 CX The functionality described on this reference page is aligned with the ISO C standard. Any  
26920 conflict between the requirements described here and the ISO C standard is unintentional. This  
26921 volume of IEEE Std 1003.1-200x defers to the ISO C standard.

26922 The *nextafter()*, *nextafterf()*, and *nextafterl()* functions shall compute the next representable  
26923 floating-point value following *x* in the direction of *y*. Thus, if *y* is less than *x*, *nextafter()* shall  
26924 return the largest representable floating-point number less than *x*. The *nextafter()*, *nextafterf()*,  
26925 and *nextafterl()* functions shall return *y* if *x* equals *y*.

26926 The *nexttoward()*, *nexttowardf()*, and *nexttowardl()* functions shall be equivalent to the |  
26927 corresponding *nextafter()* functions, except that the second parameter shall have type **long** |  
26928 **double** and the functions shall return *y* converted to the type of the function if *x* equals *y*. |

26929 An application wishing to check for error situations should set *errno* to zero and call  
26930 *feclearexcept*(FE\_ALL\_EXCEPT) before calling these functions. On return, if *errno* is non-zero or  
26931 *fetestexcept*(FE\_INVALID | FE\_DIVBYZERO | FE\_OVERFLOW | FE\_UNDERFLOW) is non-  
26932 zero, an error has occurred.

## 26933 RETURN VALUE

26934 Upon successful completion, these functions shall return the next representable floating-point  
26935 value following *x* in the direction of *y*.

26936 If *x*==*y*, *y* (of the type *x*) shall be returned.

26937 If *x* is finite and the correct function value would overflow, a range error shall occur and  
26938 ±HUGE\_VAL, ±HUGE\_VALF, and ±HUGE\_VALL (with the same sign as *x*) shall be returned as  
26939 appropriate for the return type of the function.

26940 MX If *x* or *y* is NaN, a NaN shall be returned.

26941 If *x*!=*y* and the correct function value is subnormal, zero, or underflows, a range error shall  
26942 occur, and either the correct function value (if representable) or 0.0 shall be returned.

## 26943 ERRORS

26944 These functions shall fail if:

26945 Range Error The correct value overflows

26946 If the integer expression (math\_errhandling & MATH\_ERRNO) is non-zero, |  
26947 then *errno* shall be set to [ERANGE]. If the integer expression |  
26948 (math\_errhandling & MATH\_ERREXCEPT) is non-zero, then the overflow |  
26949 floating-point exception shall be raised. |

26950 MX Range Error The correct value is subnormal or underflows

26951 If the integer expression (`math_errhandling & MATH_ERRNO`) is non-zero, |  
26952 then *errno* shall be set to [ERANGE]. If the integer expression |  
26953 (`math_errhandling & MATH_ERREXCEPT`) is non-zero, then the underflow |  
26954 floating-point exception shall be raised. |

**26955 EXAMPLES**

26956 None.

**26957 APPLICATION USAGE**

26958 On error, the expressions (`math_errhandling & MATH_ERRNO`) and (`math_errhandling &`  
26959 `MATH_ERREXCEPT`) are independent of each other, but at least one of them must be non-zero.

**26960 RATIONALE**

26961 None.

**26962 FUTURE DIRECTIONS**

26963 None.

**26964 SEE ALSO**

26965 *feclearexcept()*, *fetestexcept()*, the Base Definitions volume of IEEE Std 1003.1-200x, Section 4.18, |  
26966 Treatment of Error Conditions for Mathematical Functions, <**math.h**> |

**26967 CHANGE HISTORY**

26968 First released in Issue 4, Version 2.

**26969 Issue 5**

26970 Moved from X/OPEN UNIX extension to BASE.

**26971 Issue 6**

26972 The *nextafter()* function is no longer marked as an extension.

26973 The *nextafterf()*, *nextafterl()*, *nexttoward()*, *nexttowardf()*, *nexttowardl()* functions are added for  
26974 alignment with the ISO/IEC 9899:1999 standard.

26975 The DESCRIPTION, RETURN VALUE, ERRORS, and APPLICATION USAGE sections are  
26976 revised to align with the ISO/IEC 9899:1999 standard.

26977 IEC 60559:1989 standard floating-point extensions over the ISO/IEC 9899:1999 standard are  
26978 marked.

26979 **NAME**

26980           nexttoward, nexttowardf, nexttowardl — next representable floating-point number

26981 **SYNOPSIS**

26982           #include &lt;math.h&gt;

26983           double nexttoward(double *x*, long double *y*);26984           float nexttowardf(float *x*, long double *y*);26985           long double nexttowardl(long double *x*, long double *y*);26986 **DESCRIPTION**26987           Refer to *nextafter()*.



## 26988 NAME

26989 nftw — walk a file tree

## 26990 SYNOPSIS

26991 xSI #include &lt;ftw.h&gt;

```
26992 int nftw(const char *path, int (*fn)(const char *,
26993     const struct stat *, int, struct FTW *), int depth, int flags);
26994
```

## 26995 DESCRIPTION

26996 The *nftw()* function shall recursively descend the directory hierarchy rooted in *path*. The *nftw()*  
 26997 function has a similar effect to *ftw()* except that it takes an additional argument *flags*, which is a  
 26998 bitwise-inclusive OR of zero or more of the following flags:

26999 FTW\_CHDIR If set, *nftw()* shall change the current working directory to each directory as it  
 27000 reports files in that directory. If clear, *nftw()* shall not change the current  
 27001 working directory.

27002 FTW\_DEPTH If set, *nftw()* shall report all files in a directory before reporting the directory  
 27003 itself. If clear, *nftw()* shall report any directory before reporting the files in that  
 27004 directory.

27005 FTW\_MOUNT If set, *nftw()* shall only report files in the same file system as *path*. If clear,  
 27006 *nftw()* shall report all files encountered during the walk.

27007 FTW\_PHYS If set, *nftw()* shall perform a physical walk and shall not follow symbolic links.

27008 If FTW\_PHYS is clear and FTW\_DEPTH is set, *nftw()* shall follow links instead of reporting  
 27009 them, but shall not report any directory that would be a descendant of itself. If FTW\_PHYS is  
 27010 clear and FTW\_DEPTH is clear, *nftw()* shall follow links instead of reporting them, but shall not  
 27011 report the contents of any directory that would be a descendant of itself.

27012 At each file it encounters, *nftw()* shall call the user-supplied function *fn* with four arguments:

- 27013 • The first argument is the pathname of the object.
- 27014 • The second argument is a pointer to the **stat** buffer containing information on the object.
- 27015 • The third argument is an integer giving additional information. Its value is one of the  
 27016 following:

27017 FTW\_F The object is a file.

27018 FTW\_D The object is a directory.

27019 FTW\_DP The object is a directory and subdirectories have been visited. (This condition  
 27020 shall only occur if the FTW\_DEPTH flag is included in *flags*.)

27021 FTW\_SL The object is a symbolic link. (This condition shall only occur if the FTW\_PHYS  
 27022 flag is included in *flags*.)

27023 FTW\_SLN The object is a symbolic link that does not name an existing file. (This  
 27024 condition shall only occur if the FTW\_PHYS flag is not included in *flags*.)

27025 FTW\_DNR The object is a directory that cannot be read. The *fn* function shall not be called  
 27026 for any of its descendants.

27027 FTW\_NS The *stat()* function failed on the object because of lack of appropriate  
 27028 permission. The **stat** buffer passed to *fn* is undefined. Failure of *stat()* for any  
 27029 other reason is considered an error and *nftw()* shall return  $-1$ .

27030 • The fourth argument is a pointer to an **FTW** structure. The value of **base** is the offset of the  
 27031 object's filename in the pathname passed as the first argument to *fn*. The value of **level**  
 27032 indicates depth relative to the root of the walk, where the root level is 0.

27033 The results are unspecified if the application-supplied *fn* function does not preserve the current  
 27034 working directory.

27035 The argument *depth* sets the maximum number of file descriptors that are shall be used by *nftw()*  
 27036 while traversing the file tree. At most one file descriptor shall be used for each directory level.

27037 The *nftw()* function need not be reentrant. A function that is not required to be reentrant is not  
 27038 required to be thread-safe.

#### 27039 RETURN VALUE

27040 The *nftw()* function shall continue until the first of the following conditions occurs:

- 27041 • An invocation of *fn* shall return a non-zero value, in which case *nftw()* shall return that value.
- 27042 • The *nftw()* function detects an error other than [EACCES] (see FTW\_DNR and FTW\_NS  
 27043 above), in which case *nftw()* shall return  $-1$  and set *errno* to indicate the error.
- 27044 • The tree is exhausted, in which case *nftw()* shall return 0.

#### 27045 ERRORS

27046 The *nftw()* function shall fail if:

27047 [EACCES] Search permission is denied for any component of *path* or read permission is  
 27048 denied for *path*, or *fn* returns  $-1$  and does not reset *errno*.

27049 [ELOOP] A loop exists in symbolic links encountered during resolution of the *path*  
 27050 argument.

27051 [ENAMETOOLONG]

27052 The length of the *path* argument exceeds {PATH\_MAX} or a pathname  
 27053 component is longer than {NAME\_MAX}.

27054 [ENOENT] A component of *path* does not name an existing file or *path* is an empty string.

27055 [ENOTDIR] A component of *path* is not a directory.

27056 [EOVERFLOW] A field in the **stat** structure cannot be represented correctly in the current  
 27057 programming environment for one or more files found in the file hierarchy.

27058 The *nftw()* function may fail if:

27059 [ELOOP] More than {SYMLOOP\_MAX} symbolic links were encountered during  
 27060 resolution of the *path* argument.

27061 [EMFILE] {OPEN\_MAX} file descriptors are currently open in the calling process.

27062 [ENAMETOOLONG]

27063 Pathname resolution of a symbolic link produced an intermediate result  
 27064 whose length exceeds {PATH\_MAX}.

27065 [ENFILE] Too many files are currently open in the system.

27066 In addition, *errno* may be set if the function pointed by *fn* causes *errno* to be set.

27067 **EXAMPLES**

27068 The following example walks the `/tmp` directory and its subdirectories, calling the `nftw()`  
 27069 function for every directory entry, to a maximum of 5 levels deep.

```
27070 #include <ftw.h>
27071 ...
27072 int nftwfunc(const char *, const struct stat *, int, struct FTW *);
27073
27074 int nftwfunc(const char *filename, const struct stat *statptr,
27075             int fileflags, struct FTW *pftw)
27076 {
27077     return 0;
27078 }
27079 ...
27080 char *startpath = "/tmp";
27081 int depth = 5;
27082 int flags = FTW_CHDIR | FTW_DEPTH | FTW_MOUNT;
27083 int ret;
27084
27085 ret = nftw(startpath, nftwfunc, depth, flags);
```

27084 **APPLICATION USAGE**

27085 None.

27086 **RATIONALE**

27087 None.

27088 **FUTURE DIRECTIONS**

27089 None.

27090 **SEE ALSO**

27091 `lstat()`, `opendir()`, `readdir()`, `stat()`, the Base Definitions volume of IEEE Std 1003.1-200x, `<ftw.h>`

27092 **CHANGE HISTORY**

27093 First released in Issue 4, Version 2.

27094 **Issue 5**

27095 Moved from X/OPEN UNIX extension to BASE.

27096 In the DESCRIPTION, the definition of the *depth* argument is clarified.

27097 **Issue 6**

27098 The Open Group Base Resolution bwg97-003 is applied.

27099 The ERRORS section is updated as follows: |

27100 • The wording of the mandatory [ELOOP] error condition is updated. |

27101 • A second optional [ELOOP] error condition is added. |

27102 • The [EOVERFLOW] mandatory error condition is added. |

27103 Text is added to the DESCRIPTION to say that the `nftw()` function need not be reentrant and |  
 27104 that the results are unspecified if the application-supplied *fn* function does not preserve the |  
 27105 current working directory. |

27106 **NAME**

27107 nice — change the nice value of a process

27108 **SYNOPSIS**27109 XSI `#include <unistd.h>`27110 `int nice(int incr);`

27111

27112 **DESCRIPTION**

27113 The *nice()* function shall add the value of *incr* to the nice value of the calling process. A process' |  
 27114 nice value is a non-negative number for which a more positive value shall result in less favorable |  
 27115 scheduling.

27116 A maximum nice value of  $2^{\{NZERO\}}-1$  and a minimum nice value of 0 shall be imposed by the |  
 27117 system. Requests for values above or below these limits shall result in the nice value being set to |  
 27118 the corresponding limit. Only a process with appropriate privileges can lower the nice value.

27119 PS|TPS Calling the *nice()* function has no effect on the priority of processes or threads with policy |  
 27120 SCHED\_FIFO or SCHED\_RR. The effect on processes or threads with other scheduling policies |  
 27121 is implementation-defined.

27122 The nice value set with *nice()* shall be applied to the process. If the process is multi-threaded, the |  
 27123 nice value shall affect all system scope threads in the process.

27124 As  $-1$  is a permissible return value in a successful situation, an application wishing to check for |  
 27125 error situations should set *errno* to 0, then call *nice()*, and if it returns  $-1$ , check to see whether |  
 27126 *errno* is non-zero.

27127 **RETURN VALUE**

27128 Upon successful completion, *nice()* shall return the new nice value  $-\{NZERO\}$ . Otherwise,  $-1$  |  
 27129 shall be returned, the process' nice value shall not be changed, and *errno* shall be set to indicate |  
 27130 the error.

27131 **ERRORS**27132 The *nice()* function shall fail if:

27133 [EPERM] The *incr* argument is negative and the calling process does not have |  
 27134 appropriate privileges.

27135 **EXAMPLES**27136 **Changing the Nice Value**

27137 The following example adds the value of the *incr* argument,  $-20$ , to the nice value of the calling |  
 27138 process.

27139 `#include <unistd.h>`27140 `...`27141 `int incr = -20;`27142 `int ret;`27143 `ret = nice(incr);`27144 **APPLICATION USAGE**

27145 None.

27146 **RATIONALE**

27147           None.

27148 **FUTURE DIRECTIONS**

27149           None.

27150 **SEE ALSO**27151           *getpriority()*, *setpriority()*, the Base Definitions volume of IEEE Std 1003.1-200x, <**limits.h**>,  
27152           <**unistd.h**>27153 **CHANGE HISTORY**

27154           First released in Issue 1. Derived from Issue 1 of the SVID.

27155 **Issue 5**27156           A statement is added to the description indicating the effects of this function on the different  
27157           scheduling policies and multi-threaded processes.

27158 **NAME**

27159 nl\_langinfo — language information

27160 **SYNOPSIS**

27161 xSI #include &lt;langinfo.h&gt;

27162 char \*nl\_langinfo(nl\_item item);

27163

27164 **DESCRIPTION**

27165 The *nl\_langinfo()* function shall return a pointer to a string containing information relevant to  
27166 the particular language or cultural area defined in the program's locale (see <langinfo.h>). The  
27167 manifest constant names and values of *item* are defined in <langinfo.h>. For example:

27168 nl\_langinfo(ABDAY\_1)

27169 would return a pointer to the string "Dom" if the identified language was Portuguese, and  
27170 "Sun" if the identified language was English.

27171 Calls to *setlocale()* with a category corresponding to the category of *item* (see <langinfo.h>), or to  
27172 the category *LC\_ALL*, may overwrite the array pointed to by the return value.

27173 The *nl\_langinfo()* function need not be reentrant. A function that is not required to be reentrant is  
27174 not required to be thread-safe.

27175 **RETURN VALUE**

27176 In a locale where *langinfo* data is not defined, *nl\_langinfo()* shall return a pointer to the  
27177 corresponding string in the POSIX locale. In all locales, *nl\_langinfo()* shall return a pointer to an  
27178 empty string if *item* contains an invalid setting.

27179 This pointer may point to static data that may be overwritten on the next call.

27180 **ERRORS**

27181 No errors are defined.

27182 **EXAMPLES**27183 **Getting Date and Time Formatting Information**

27184 The following example returns a pointer to a string containing date and time formatting  
27185 information, as defined in the *LC\_TIME* category of the current locale.

27186 #include &lt;time.h&gt;

27187 #include &lt;langinfo.h&gt;

27188 ...

27189 strftime(datestring, sizeof(datestring), nl\_langinfo(D\_T\_FMT), tm);

27190 ...

27191 **APPLICATION USAGE**

27192 The array pointed to by the return value should not be modified by the program, but may be  
27193 modified by further calls to *nl\_langinfo()*.

27194 **RATIONALE**

27195 None.

27196 **FUTURE DIRECTIONS**

27197 None.

27198 **SEE ALSO**

27199        *setlocale()*, the Base Definitions volume of IEEE Std 1003.1-200x, <langinfo.h>, <nl\_types.h>, the  
27200        Base Definitions volume of IEEE Std 1003.1-200x, Chapter 7, Locale

27201 **CHANGE HISTORY**

27202        First released in Issue 2.

27203 **Issue 5**

27204        The last paragraph of the DESCRIPTION is moved from the APPLICATION USAGE section.

27205        A note indicating that this function need not be reentrant is added to the DESCRIPTION.

27206 **NAME**

27207 nrand48 — generate uniformly distributed pseudo-random non-negative long integers

27208 **SYNOPSIS**

27209 xSI #include <stdlib.h>

27210 long nrand48(unsigned short xsubi[3]);

27211

27212 **DESCRIPTION**

27213 Refer to *drand48()*.



27214 **NAME**

27215           ntohl — convert values between host and network byte order

27216 **SYNOPSIS**

27217           #include <arpa/inet.h>

27218           uint32\_t ntohl(uint32\_t *netlong*);

27219 **DESCRIPTION**

27220           Refer to *htonl()*.

27221 **NAME**

27222       ntohs — convert values between host and network byte order

27223 **SYNOPSIS**

27224       #include <arpa/inet.h>

27225       uint16\_t ntohs(uint16\_t *netshort*);

27226 **DESCRIPTION**

27227       Refer to *htonl()*.

## 27228 NAME

27229 open — open a file

## 27230 SYNOPSIS

27231 OH #include &lt;sys/stat.h&gt;

27232 #include &lt;fcntl.h&gt;

27233 int open(const char \*path, int oflag, ... );

## 27234 DESCRIPTION

27235 The *open()* function shall establish the connection between a file and a file descriptor. It shall  
 27236 create an open file description that refers to a file and a file descriptor that refers to that open file  
 27237 description. The file descriptor is used by other I/O functions to refer to that file. The *path*  
 27238 argument points to a pathname naming the file.

27239 The *open()* function shall return a file descriptor for the named file that is the lowest file  
 27240 descriptor not currently open for that process. The open file description is new, and therefore the  
 27241 file descriptor shall not share it with any other process in the system. The FD\_CLOEXEC file  
 27242 descriptor flag associated with the new file descriptor shall be cleared.

27243 The file offset used to mark the current position within the file shall be set to the beginning of the  
 27244 file.

27245 The file status flags and file access modes of the open file description shall be set according to  
 27246 the value of *oflag*.

27247 Values for *oflag* are constructed by a bitwise-inclusive OR of flags from the following list,  
 27248 defined in <fcntl.h>. Applications shall specify exactly one of the first three values (file access  
 27249 modes) below in the value of *oflag*:

27250 O\_RDONLY Open for reading only.

27251 O\_WRONLY Open for writing only.

27252 O\_RDWR Open for reading and writing. The result is undefined if this flag is applied to  
27253 a FIFO.

27254 Any combination of the following may be used:

27255 O\_APPEND If set, the file offset shall be set to the end of the file prior to each write.

27256 O\_CREAT If the file exists, this flag has no effect except as noted under O\_EXCL below. |  
 27257 Otherwise, the file shall be created; the user ID of the file shall be set to the |  
 27258 effective user ID of the process; the group ID of the file shall be set to the |  
 27259 group ID of the file's parent directory or to the effective group ID of the |  
 27260 process; and the access permission bits (see <sys/stat.h>) of the file mode shall |  
 27261 be set to the value of the third argument taken as type **mode\_t** modified as |  
 27262 follows: a bitwise AND is performed on the file-mode bits and the |  
 27263 corresponding bits in the complement of the process' file mode creation mask. |  
 27264 Thus, all bits in the file mode whose corresponding bit in the file mode |  
 27265 creation mask is set are cleared. When bits other than the file permission bits |  
 27266 are set, the effect is unspecified. The third argument does not affect whether |  
 27267 the file is open for reading, writing, or for both. Implementations shall provide |  
 27268 a way to initialize the file's group ID to the group ID of the parent directory. |  
 27269 Implementations may, but need not, provide an implementation-defined way |  
 27270 to initialize the file's group ID to the effective group ID of the calling process. |

27271 SIO O\_DSYNC Write I/O operations on the file descriptor shall complete as defined by |  
 27272 synchronized I/O data integrity completion. |

27273	O_EXCL	If O_CREAT and O_EXCL are set, <i>open()</i> shall fail if the file exists. The check for the existence of the file and the creation of the file if it does not exist shall be atomic with respect to other threads executing <i>open()</i> naming the same filename in the same directory with O_EXCL and O_CREAT set. If O_EXCL and O_CREAT are set, and <i>path</i> names a symbolic link, <i>open()</i> shall fail and set <i>errno</i> to [EEXIST], regardless of the contents of the symbolic link. If O_EXCL is set and O_CREAT is not set, the result is undefined.
27274		
27275		
27276		
27277		
27278		
27279		
27280	O_NOCTTY	If set and <i>path</i> identifies a terminal device, <i>open()</i> shall not cause the terminal device to become the controlling terminal for the process.
27281		
27282	O_NONBLOCK	When opening a FIFO with O_RDONLY or O_WRONLY set:
27283		<ul style="list-style-type: none"> <li>• If O_NONBLOCK is set, an <i>open()</i> for reading-only shall return without delay. An <i>open()</i> for writing-only shall return an error if no process currently has the file open for reading.</li> </ul>
27284		
27285		
27286		<ul style="list-style-type: none"> <li>• If O_NONBLOCK is clear, an <i>open()</i> for reading-only shall block the calling thread until a thread opens the file for writing. An <i>open()</i> for writing-only shall block the calling thread until a thread opens the file for reading.</li> </ul>
27287		
27288		
27289		
27290		When opening a block special or character special file that supports non-blocking opens:
27291		
27292		<ul style="list-style-type: none"> <li>• If O_NONBLOCK is set, the <i>open()</i> function shall return without blocking for the device to be ready or available. Subsequent behavior of the device is device-specific.</li> </ul>
27293		
27294		
27295		<ul style="list-style-type: none"> <li>• If O_NONBLOCK is clear, the <i>open()</i> function shall block the calling thread until the device is ready or available before returning.</li> </ul>
27296		
27297		Otherwise, the behavior of O_NONBLOCK is unspecified.
27298	SIO O_RSYNC	Read I/O operations on the file descriptor shall complete at the same level of integrity as specified by the O_DSYNC and O_SYNC flags. If both O_DSYNC and O_RSYNC are set in <i>oflag</i> , all I/O operations on the file descriptor shall complete as defined by synchronized I/O data integrity completion. If both O_SYNC and O_RSYNC are set in flags, all I/O operations on the file descriptor shall complete as defined by synchronized I/O file integrity completion.
27299		
27300		
27301		
27302		
27303		
27304		
27305	SIO O_SYNC	Write I/O operations on the file descriptor shall complete as defined by synchronized I/O file integrity completion.
27306		
27307	O_TRUNC	If the file exists and is a regular file, and the file is successfully opened O_RDWR or O_WRONLY, its length shall be truncated to 0, and the mode and owner shall be unchanged. It shall have no effect on FIFO special files or terminal device files. Its effect on other file types is implementation-defined. The result of using O_TRUNC with O_RDONLY is undefined.
27308		
27309		
27310		
27311		
27312		If O_CREAT is set and the file did not previously exist, upon successful completion, <i>open()</i> shall 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 <i>st_mtime</i> fields of the parent directory.
27313		
27314		
27315		If O_TRUNC is set and the file did previously exist, upon successful completion, <i>open()</i> shall mark for update the <i>st_ctime</i> and <i>st_mtime</i> fields of the file.
27316		

27317 SIO 27318	If both the O_SYNC and O_DSYNC flags are set, the effect is as if only the O_SYNC flag was set.	
27319 XSR 27320 27321 27322 27323 27324	If <i>path</i> refers to a STREAMS file, <i>oflag</i> may be constructed from O_NONBLOCK OR'ed with either O_RDONLY, O_WRONLY, or O_RDWR. Other flag values are not applicable to STREAMS devices and shall have no effect on them. The value O_NONBLOCK affects the operation of STREAMS drivers and certain functions applied to file descriptors associated with STREAMS files. For STREAMS drivers, the implementation of O_NONBLOCK is device-specific.	
27325 XSI 27326 27327	If <i>path</i> names the master side of a pseudo-terminal device, then it is unspecified whether <i>open()</i> locks the slave side so that it cannot be opened. Conforming applications shall call <i>unlockpt()</i> before opening the slave side.	
27328 27329	The largest value that can be represented correctly in an object of type <b>off_t</b> shall be established as the offset maximum in the open file description.	
27330	<b>RETURN VALUE</b>	
27331 27332 27333	Upon successful completion, the function shall open the file and return a non-negative integer representing the lowest numbered unused file descriptor. Otherwise, <code>-1</code> shall be returned and <i>errno</i> set to indicate the error. No files shall be created or modified if the function returns <code>-1</code> .	
27334	<b>ERRORS</b>	
27335	The <i>open()</i> function shall fail if:	
27336 27337 27338 27339	[EACCES]	Search permission is denied on a component of the path prefix, or the file exists and the permissions specified by <i>oflag</i> are denied, or the file does not exist and write permission is denied for the parent directory of the file to be created, or O_TRUNC is specified and write permission is denied.
27340	[EEXIST]	O_CREAT and O_EXCL are set, and the named file exists.
27341	[EINTR]	A signal was caught during <i>open()</i> .
27342 SIO	[EINVAL]	The implementation does not support synchronized I/O for this file.
27343 XSR 27344	[EIO]	The <i>path</i> argument names a STREAMS file and a hangup or error occurred during the <i>open()</i> .
27345	[EISDIR]	The named file is a directory and <i>oflag</i> includes O_WRONLY or O_RDWR.
27346 27347	[ELOOP]	A loop exists in symbolic links encountered during resolution of the <i>path</i> argument.
27348	[EMFILE]	{OPEN_MAX} file descriptors are currently open in the calling process.
27349	[ENAMETOOLONG]	
27350 27351		The length of the <i>path</i> argument exceeds {PATH_MAX} or a pathname component is longer than {NAME_MAX}.
27352	[ENFILE]	The maximum allowable number of files is currently open in the system.
27353 27354 27355	[ENOENT]	O_CREAT is not set and the named file does not exist; or O_CREAT is set and either the path prefix does not exist or the <i>path</i> argument points to an empty string.
27356 XSR 27357	[ENOSR]	The <i>path</i> argument names a STREAMS-based file and the system is unable to allocate a STREAM.
27358 27359	[ENOSPC]	The directory or file system that would contain the new file cannot be expanded, the file does not exist, and O_CREAT is specified.

27360	[ENOTDIR]	A component of the path prefix is not a directory.
27361	[ENXIO]	O_NONBLOCK is set, the named file is a FIFO, O_WRONLY is set, and no process has the file open for reading.
27362		
27363	[ENXIO]	The named file is a character special or block special file, and the device associated with this special file does not exist.
27364		
27365	[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> .
27366		
27367	[EROFS]	The named file resides on a read-only file system and either O_WRONLY, O_RDWR, O_CREAT (if file does not exist), or O_TRUNC is set in the <i>oflag</i> argument.
27368		
27369		
27370		The <i>open()</i> function may fail if:
27371 XSI	[EAGAIN]	The <i>path</i> argument names the slave side of a pseudo-terminal device that is locked.
27372		
27373	[EINVAL]	The value of the <i>oflag</i> argument is not valid.
27374	[ELOOP]	More than {SYMLOOP_MAX} symbolic links were encountered during resolution of the <i>path</i> argument.
27375		
27376	[ENAMETOOLONG]	
27377		As a result of encountering a symbolic link in resolution of the <i>path</i> argument, the length of the substituted pathname string exceeded {PATH_MAX}.
27378		
27379 XSR	[ENOMEM]	The <i>path</i> argument names a STREAMS file and the system is unable to allocate resources.
27380		
27381	[ETXTBSY]	The file is a pure procedure (shared text) file that is being executed and <i>oflag</i> is O_WRONLY or O_RDWR.
27382		

## 27383 EXAMPLES

### 27384 Opening a File for Writing by the Owner

27385 The following example opens the file `/tmp/file`, either by creating it (if it does not already exist),  
 27386 or by truncating its length to 0 (if it does exist). In the former case, if the call creates a new file,  
 27387 the access permission bits in the file mode of the file are set to permit reading and writing by the  
 27388 owner, and to permit reading only by group members and others.

27389 If the call to *open()* is successful, the file is opened for writing.

```

27390 #include <fcntl.h>
27391 ...
27392 int fd;
27393 mode_t mode = S_IRUSR | S_IWUSR | S_IRGRP | S_IROTH;
27394 char *filename = "/tmp/file";
27395 ...
27396 fd = open(filename, O_WRONLY | O_CREAT | O_TRUNC, mode);
27397 ...

```

27398 **Opening a File Using an Existence Check**

27399 The following example uses the *open()* function to try to create the **LOCKFILE** file and open it |  
 27400 for writing. Since the *open()* function specifies the **O\_EXCL** flag, the call fails if the file already |  
 27401 exists. In that case, the program assumes that someone else is updating the password file and  
 27402 exits.

```
27403 #include <fcntl.h>
27404 #include <stdio.h>
27405 #include <stdlib.h>

27406 #define LOCKFILE "/etc/ptmp"
27407 ...
27408 int pfd; /* Integer for file descriptor returned by open() call. */
27409 ...
27410 if ((pfd = open(LOCKFILE, O_WRONLY | O_CREAT | O_EXCL,
27411             S_IRUSR | S_IWUSR | S_IRGRP | S_IROTH)) == -1)
27412 {
27413     fprintf(stderr, "Cannot open /etc/ptmp. Try again later.\n");
27414     exit(1);
27415 }
27416 ...
```

27417 **Opening a File for Writing**

27418 The following example opens a file for writing, creating the file if it does not already exist. If the  
 27419 file does exist, the system truncates the file to zero bytes.

```
27420 #include <fcntl.h>
27421 #include <stdio.h>
27422 #include <stdlib.h>

27423 #define LOCKFILE "/etc/ptmp"
27424 ...
27425 int pfd;
27426 char filename[PATH_MAX+1];
27427 ...
27428 if ((pfd = open(filename, O_WRONLY | O_CREAT | O_TRUNC,
27429             S_IRUSR | S_IWUSR | S_IRGRP | S_IROTH)) == -1)
27430 {
27431     perror("Cannot open output file\n"); exit(1);
27432 }
27433 ...
```

27434 **APPLICATION USAGE**

27435 None.

27436 **RATIONALE**

27437 Except as specified in this volume of IEEE Std 1003.1-200x, the flags allowed in *oflag* are not  
 27438 mutually-exclusive and any number of them may be used simultaneously.

27439 Some implementations permit opening FIFOs with **O\_RDWR**. Since FIFOs could be  
 27440 implemented in other ways, and since two file descriptors can be used to the same effect, this  
 27441 possibility is left as undefined.

27442 See *getgroups()* about the group of a newly created file.

27443 The use of *open()* to create a regular file is preferable to the use of *creat()*, because the latter is  
27444 redundant and included only for historical reasons.

27445 The use of the O\_TRUNC flag on FIFOs and directories (pipes cannot be *open()*-ed) must be  
27446 permissible without unexpected side effects (for example, *creat()* on a FIFO must not remove  
27447 data). Since terminal special files might have type-ahead data stored in the buffer, O\_TRUNC  
27448 should not affect their content, particularly if a program that normally opens a regular file  
27449 should open the current controlling terminal instead. Other file types, particularly  
27450 implementation-defined ones, are left implementation-defined.

27451 IEEE Std 1003.1-200x permits [EACCES] to be returned for conditions other than those explicitly  
27452 listed.

27453 The O\_NOCTTY flag was added to allow applications to avoid unintentionally acquiring a  
27454 controlling terminal as a side effect of opening a terminal file. This volume of  
27455 IEEE Std 1003.1-200x does not specify how a controlling terminal is acquired, but it allows an  
27456 implementation to provide this on *open()* if the O\_NOCTTY flag is not set and other conditions  
27457 specified in the Base Definitions volume of IEEE Std 1003.1-200x, Chapter 11, General Terminal  
27458 Interface are met. The O\_NOCTTY flag is an effective no-op if the file being opened is not a  
27459 terminal device.

27460 In historical implementations the value of O\_RDONLY is zero. Because of that, it is not possible  
27461 to detect the presence of O\_RDONLY and another option. Future implementations should  
27462 encode O\_RDONLY and O\_WRONLY as bit flags so that:

```
27463 O_RDONLY | O_WRONLY == O_RDWR
```

27464 In general, the *open()* function follows the symbolic link if *path* names a symbolic link. However,  
27465 the *open()* function, when called with O\_CREAT and O\_EXCL, is required to fail with [EEXIST]  
27466 if *path* names an existing symbolic link, even if the symbolic link refers to a nonexistent file. This  
27467 behavior is required so that privileged applications can create a new file in a known location  
27468 without the possibility that a symbolic link might cause the file to be created in a different  
27469 location.

27470 For example, a privileged application that must create a file with a predictable name in a user-  
27471 writable directory, such as the user's home directory, could be compromised if the user creates a  
27472 symbolic link with that name that refers to a nonexistent file in a system directory. If the user can  
27473 influence the contents of a file, the user could compromise the system by creating a new system  
27474 configuration or spool file that would then be interpreted by the system. The test for a symbolic  
27475 link which refers to a nonexistent file must be atomic with the creation of a new file.

27476 The POSIX.1-1990 standard required that the group ID of a newly created file be set to the group  
27477 ID of its parent directory or to the effective group ID of the creating process. FIPS 151-2 required  
27478 that implementations provide a way to have the group ID be set to the group ID of the  
27479 containing directory, but did not prohibit implementations also supporting a way to set the  
27480 group ID to the effective group ID of the creating process. Conforming applications should not  
27481 assume which group ID will be used. If it matters, an application can use *chown()* to set the  
27482 group ID after the file is created, or determine under what conditions the implementation will  
27483 set the desired group ID.

#### 27484 FUTURE DIRECTIONS

27485 None.

#### 27486 SEE ALSO

27487 *chmod()*, *close()*, *creat()*, *dup()*, *fcntl()*, *lseek()*, *read()*, *umask()*, *unlockpt()*, *write()*, the Base  
27488 Definitions volume of IEEE Std 1003.1-200x, <fcntl.h>, <sys/stat.h>, <sys/types.h>



27489 **CHANGE HISTORY**

27490 First released in Issue 1. Derived from Issue 1 of the SVID.

27491 **Issue 5**

27492 The DESCRIPTION is updated for alignment with the POSIX Realtime Extension and the POSIX  
27493 Threads Extension.

27494 Large File Summit extensions are added.

27495 **Issue 6**

27496 In the SYNOPSIS, the optional include of the `<sys/types.h>` header is removed.

27497 The following new requirements on POSIX implementations derive from alignment with the  
27498 Single UNIX Specification:

27499 • The requirement to include `<sys/types.h>` has been removed. Although `<sys/types.h>` was  
27500 required for conforming implementations of previous POSIX specifications, it was not  
27501 required for UNIX applications.

27502 • In the DESCRIPTION, `O_CREAT` is amended to state that the group ID of the file is set to the  
27503 group ID of the file's parent directory or to the effective group ID of the process. This is a  
27504 FIPS requirement.

27505 • In the DESCRIPTION, text is added to indicate setting of the offset maximum in the open file  
27506 description. This change is to support large files.

27507 • In the ERRORS section, the `[E_OVERFLOW]` condition is added. This change is to support  
27508 large files.

27509 • The `[ENXIO]` mandatory error condition is added.

27510 • The `[EINVAL]`, `[ENAMETOOLONG]`, and `[ETXTBSY]` optional error conditions are added.

27511 The DESCRIPTION and ERRORS sections are updated so that items related to the optional XSI  
27512 STREAMS Option Group are marked.

27513 The following changes were made to align with the IEEE P1003.1a draft standard:

27514 • An explanation is added of the effect of the `O_CREAT` and `O_EXCL` flags when the path  
27515 refers to a symbolic link.

27516 • The `[ELOOP]` optional error condition is added.

27517 The DESCRIPTION is updated to avoid use of the term “must” for application requirements.

27518 The DESCRIPTION of `O_EXCL` is updated in response to IEEE PASC Interpretation 1003.1c #48.

27519 **NAME**

27520            opendir — open a directory

27521 **SYNOPSIS**

27522            #include <dirent.h>

27523            DIR \*opendir(const char \*dirname);

27524 **DESCRIPTION**

27525            The *opendir()* function shall open a directory stream corresponding to the directory named by  
 27526            the *dirname* argument. The directory stream is positioned at the first entry. If the type **DIR** is  
 27527            implemented using a file descriptor, applications shall only be able to open up to a total of  
 27528            {OPEN\_MAX} files and directories.

27529 **RETURN VALUE**

27530            Upon successful completion, *opendir()* shall return a pointer to an object of type **DIR**.  
 27531            Otherwise, a null pointer shall be returned and *errno* set to indicate the error.

27532 **ERRORS**

27533            The *opendir()* function shall fail if:

27534            [EACCES]            Search permission is denied for the component of the path prefix of *dirname* or  
 27535            read permission is denied for *dirname*.

27536            [ELOOP]            A loop exists in symbolic links encountered during resolution of the *dirname*  
 27537            argument.

27538            [ENAMETOOLONG]       The length of the *dirname* argument exceeds {PATH\_MAX} or a pathname  
 27539            component is longer than {NAME\_MAX}.  
 27540

27541            [ENOENT]            A component of *dirname* does not name an existing directory or *dirname* is an  
 27542            empty string.

27543            [ENOTDIR]            A component of *dirname* is not a directory.

27544            The *opendir()* function may fail if:

27545            [ELOOP]            More than {SYMLOOP\_MAX} symbolic links were encountered during  
 27546            resolution of the *dirname* argument.

27547            [EMFILE]            {OPEN\_MAX} file descriptors are currently open in the calling process.

27548            [ENAMETOOLONG]       As a result of encountering a symbolic link in resolution of the *dirname*  
 27549            argument, the length of the substituted pathname string exceeded  
 27550            {PATH\_MAX}.  
 27551

27552            [ENFILE]            Too many files are currently open in the system.

27553 **EXAMPLES**27554 **Open a Directory Stream**

27555 The following program fragment demonstrates how the *opendir()* function is used.

```

27556 #include <sys/types.h>
27557 #include <dirent.h>
27558 #include <libgen.h>
27559 ...
27560     DIR *dir;
27561     struct dirent *dp;
27562 ...
27563     if ((dir = opendir(".")) == NULL) {
27564         perror("Cannot open .");
27565         exit(1);
27566     }
27567     while ((dp = readdir(dir)) != NULL) {
27568         ...

```

27569 **APPLICATION USAGE**

27570 The *opendir()* function should be used in conjunction with *readdir()*, *closedir()*, and *rewinddir()* to  
 27571 examine the contents of the directory (see the EXAMPLES section in *readdir()*). This method is  
 27572 recommended for portability.

27573 **RATIONALE**

27574 Based on historical implementations, the rules about file descriptors apply to directory streams  
 27575 as well. However, this volume of IEEE Std 1003.1-200x does not mandate that the directory  
 27576 stream be implemented using file descriptors. The description of *closedir()* clarifies that if a file  
 27577 descriptor is used for the directory stream, it is mandatory that *closedir()* deallocate the file  
 27578 descriptor. When a file descriptor is used to implement the directory stream, it behaves as if the  
 27579 FD\_CLOEXEC had been set for the file descriptor.

27580 The directory entries for dot and dot-dot are optional. This volume of IEEE Std 1003.1-200x does  
 27581 not provide a way to test *a priori* for their existence because an application that is portable must  
 27582 be written to look for (and usually ignore) those entries. Writing code that presumes that they  
 27583 are the first two entries does not always work, as many implementations permit them to be  
 27584 other than the first two entries, with a “normal” entry preceding them. There is negligible value  
 27585 in providing a way to determine what the implementation does because the code to deal with  
 27586 dot and dot-dot must be written in any case and because such a flag would add to the list of  
 27587 those flags (which has proven in itself to be objectionable) and might be abused.

27588 Since the structure and buffer allocation, if any, for directory operations are defined by the  
 27589 implementation, this volume of IEEE Std 1003.1-200x imposes no portability requirements for  
 27590 erroneous program constructs, erroneous data, or the use of unspecified values such as the use  
 27591 or referencing of a *dirp* value or a **dirent** structure value after a directory stream has been closed  
 27592 or after a *fork()* or one of the *exec* function calls.

27593 **FUTURE DIRECTIONS**

27594 None.

27595 **SEE ALSO**

27596 *closedir()*, *lstat()*, *readdir()*, *rewinddir()*, *symlink()*, the Base Definitions volume of  
 27597 IEEE Std 1003.1-200x, **<dirent.h>**, **<limits.h>**, **<sys/types.h>**

27598 **CHANGE HISTORY**

27599 First released in Issue 2.

27600 **Issue 6**

27601 In the SYNOPSIS, the optional include of the `<sys/types.h>` header is removed.

27602 The following new requirements on POSIX implementations derive from alignment with the |  
27603 Single UNIX Specification:

27604 • The requirement to include `<sys/types.h>` has been removed. Although `<sys/types.h>` was  
27605 required for conforming implementations of previous POSIX specifications, it was not  
27606 required for UNIX applications.

27607 • The [ELOOP] mandatory error condition is added.

27608 • A second [ENAMETOOLONG] is added as an optional error condition.

27609 The following changes were made to align with the IEEE P1003.1a draft standard:

27610 • The [ELOOP] optional error condition is added.

27611 **NAME**

27612           openlog — open a connection to the logging facility

27613 **SYNOPSIS**

27614 xSI       #include &lt;syslog.h&gt;

27615           void openlog(const char \*ident, int logopt, int facility);

27616

27617 **DESCRIPTION**27618           Refer to *closelog()*.

27619 **NAME**

27620           optarg, opterr, optind, optopt — options parsing variables

27621 **SYNOPSIS**

27622           #include &lt;unistd.h&gt;

27623           extern char \*optarg;

27624           extern int opterr, optind, optopt;

27625 **DESCRIPTION**27626           Refer to *getopt()*.

27627 **NAME**

27628 pathconf — get configurable pathname variables |

27629 **SYNOPSIS**

27630 #include &lt;unistd.h&gt;

27631 long pathconf(const char \*path, int name);

27632 **DESCRIPTION**27633 Refer to *fpathconf()*.

27634 **NAME**

27635 pause — suspend the thread until a signal is received

27636 **SYNOPSIS**

27637 #include <unistd.h>

27638 int pause(void);

27639 **DESCRIPTION**

27640 The *pause()* function shall suspend the calling thread until delivery of a signal whose action is  
27641 either to execute a signal-catching function or to terminate the process.

27642 If the action is to terminate the process, *pause()* shall not return.

27643 If the action is to execute a signal-catching function, *pause()* shall return after the signal-catching  
27644 function returns.

27645 **RETURN VALUE**

27646 Since *pause()* suspends thread execution indefinitely unless interrupted by a signal, there is no  
27647 successful completion return value. A value of  $-1$  shall be returned and *errno* set to indicate the  
27648 error.

27649 **ERRORS**

27650 The *pause()* function shall fail if:

27651 [EINTR] A signal is caught by the calling process and control is returned from the  
27652 signal-catching function.

27653 **EXAMPLES**

27654 None.

27655 **APPLICATION USAGE**

27656 Many common uses of *pause()* have timing windows. The scenario involves checking a  
27657 condition related to a signal and, if the signal has not occurred, calling *pause()*. When the signal  
27658 occurs between the check and the call to *pause()*, the process often blocks indefinitely. The  
27659 *sigprocmask()* and *sigsuspend()* functions can be used to avoid this type of problem.

27660 **RATIONALE**

27661 None.

27662 **FUTURE DIRECTIONS**

27663 None.

27664 **SEE ALSO**

27665 *sigsuspend()*, the Base Definitions volume of IEEE Std 1003.1-200x, <unistd.h>

27666 **CHANGE HISTORY**

27667 First released in Issue 1. Derived from Issue 1 of the SVID.

27668 **Issue 5**

27669 The DESCRIPTION is updated for alignment with the POSIX Threads Extension.

27670 **Issue 6**

27671 The APPLICATION USAGE section is added.



27672 **NAME**

27673 pclose — close a pipe stream to or from a process

27674 **SYNOPSIS**

27675 cx #include &lt;stdio.h&gt;

27676 int pclose(FILE \*stream);

27677

27678 **DESCRIPTION**

27679 The *pclose()* function shall close a stream that was opened by *popen()*, wait for the command to  
 27680 terminate, and return the termination status of the process that was running the command  
 27681 language interpreter. However, if a call caused the termination status to be unavailable to  
 27682 *pclose()*, then *pclose()* shall return  $-1$  with *errno* set to [ECHILD] to report this situation. This can  
 27683 happen if the application calls one of the following functions:

- 27684 • *wait()*
- 27685 • *waitpid()* with a *pid* argument less than or equal to 0 or equal to the process ID of the  
 27686 command line interpreter
- 27687 • Any other function not defined in this volume of IEEE Std 1003.1-200x that could do one of  
 27688 the above

27689 In any case, *pclose()* shall not return before the child process created by *popen()* has terminated.

27690 If the command language interpreter cannot be executed, the child termination status returned  
 27691 by *pclose()* shall be as if the command language interpreter terminated using *exit(127)* or  
 27692 *\_exit(127)*.

27693 The *pclose()* function shall not affect the termination status of any child of the calling process  
 27694 other than the one created by *popen()* for the associated stream.

27695 If the argument *stream* to *pclose()* is not a pointer to a stream created by *popen()*, the result of  
 27696 *pclose()* is undefined.

27697 **RETURN VALUE**

27698 Upon successful return, *pclose()* shall return the termination status of the command language  
 27699 interpreter. Otherwise, *pclose()* shall return  $-1$  and set *errno* to indicate the error.

27700 **ERRORS**

27701 The *pclose()* function shall fail if:

27702 [ECHILD] The status of the child process could not be obtained, as described above.

27703 **EXAMPLES**

27704 None.

27705 **APPLICATION USAGE**

27706 None.

27707 **RATIONALE**

27708 There is a requirement that *pclose()* not return before the child process terminates. This is  
 27709 intended to disallow implementations that return [EINTR] if a signal is received while waiting.  
 27710 If *pclose()* returned before the child terminated, there would be no way for the application to  
 27711 discover which child used to be associated with the stream, and it could not do the cleanup  
 27712 itself.

27713 If the stream pointed to by *stream* was not created by *popen()*, historical implementations of  
 27714 *pclose()* return  $-1$  without setting *errno*. To avoid requiring *pclose()* to set *errno* in this case,  
 27715 IEEE Std 1003.1-200x makes the behavior unspecified. An application should not use *pclose()* to

27716 close any stream that was not created by *popen()*.

27717 Some historical implementations of *pclose()* either block or ignore the signals SIGINT, SIGQUIT,  
27718 and SIGHUP while waiting for the child process to terminate. Since this behavior is not  
27719 described for the *pclose()* function in IEEE Std 1003.1-200x, such implementations are not  
27720 conforming. Also, some historical implementations return [EINTR] if a signal is received, even  
27721 though the child process has not terminated. Such implementations are also considered non-  
27722 conforming.

27723 Consider, for example, an application that uses:

```
27724 popen("command", "r")
```

27725 to start *command*, which is part of the same application. The parent writes a prompt to its  
27726 standard output (presumably the terminal) and then reads from the stream. The child reads the  
27727 response from the user, does some transformation on the response (pathname expansion, |  
27728 perhaps) and writes the result to its standard output. The parent process reads the result from |  
27729 the pipe, does something with it, and prints another prompt. The cycle repeats. Assuming that  
27730 both processes do appropriate buffer flushing, this would be expected to work.

27731 To conform to IEEE Std 1003.1-200x, *pclose()* must use *waitpid()*, or some similar function,  
27732 instead of *wait()*.

27733 The code sample below illustrates how the *pclose()* function might be implemented on a system  
27734 conforming to IEEE Std 1003.1-200x.

```
27735 int pclose(FILE *stream)
27736 {
27737     int stat;
27738     pid_t pid;
27739
27740     pid = <pid for process created for stream by popen(>
27741     (void) fclose(stream);
27742     while (waitpid(pid, &stat, 0) == -1) {
27743         if (errno != EINTR){
27744             stat = -1;
27745             break;
27746         }
27747     }
27748     return(stat);
27749 }
```

#### 27749 FUTURE DIRECTIONS

27750 None.

#### 27751 SEE ALSO

27752 *fork()*, *popen()*, *waitpid()*, the Base Definitions volume of IEEE Std 1003.1-200x, <stdio.h>

#### 27753 CHANGE HISTORY

27754 First released in Issue 1. Derived from Issue 1 of the SVID.

27755 **NAME**

27756           perror — write error messages to standard error

27757 **SYNOPSIS**

27758           #include &lt;stdio.h&gt;

27759           void perror(const char \*s);

27760 **DESCRIPTION**

27761 cx       The functionality described on this reference page is aligned with the ISO C standard. Any  
 27762       conflict between the requirements described here and the ISO C standard is unintentional. This  
 27763       volume of IEEE Std 1003.1-200x defers to the ISO C standard.

27764       The *perror()* function shall map the error number accessed through the symbol *errno* to a  
 27765       language-dependent error message, which shall be written to the standard error stream as  
 27766       follows:

- 27767           • First (if *s* is not a null pointer and the character pointed to by *s* is not the null byte), the string  
 27768           pointed to by *s* followed by a colon and a <space>.
- 27769           • Then an error message string followed by a <newline>.

27770       The contents of the error message strings shall be the same as those returned by *strerror()* with  
 27771       argument *errno*.

27772 cx       The *perror()* function shall mark the file associated with the standard error stream as having  
 27773       been written (*st\_ctime*, *st\_mtime* marked for update) at some time between its successful  
 27774       completion and *exit()*, *abort()*, or the completion of *fflush()* or *fclose()* on *stderr*.

27775       The *perror()* function shall not change the orientation of the standard error stream.

27776 **RETURN VALUE**27777       The *perror()* function shall not return a value.27778 **ERRORS**

27779       No errors are defined.

27780 **EXAMPLES**27781       **Printing an Error Message for a Function**

27782       The following example replaces *bufptr* with a buffer that is the necessary size. If an error occurs,  
 27783       the *perror()* function prints a message and the program exits.

```

27784       #include <stdio.h>
27785       #include <stdlib.h>
27786       ...
27787       char *bufptr;
27788       size_t szbuf;
27789       ...
27790       if ((bufptr = malloc(szbuf)) == NULL) {
27791           perror("malloc"); exit(2);
27792       }
27793       ...
```

27794 **APPLICATION USAGE**

27795       None.

27796 **RATIONALE**

27797           None.

27798 **FUTURE DIRECTIONS**

27799           None.

27800 **SEE ALSO**27801           *strerror()*, the Base Definitions volume of IEEE Std 1003.1-200x, <stdio.h>27802 **CHANGE HISTORY**

27803           First released in Issue 1. Derived from Issue 1 of the SVID.

27804 **Issue 5**27805           A paragraph is added to the DESCRIPTION indicating that *perror()* does not change the  
27806           orientation of the standard error stream.27807 **Issue 6**

27808           Extensions beyond the ISO C standard are now marked.

27809 **NAME**

27810 pipe — create an interprocess channel

27811 **SYNOPSIS**

27812 #include &lt;unistd.h&gt;

27813 int pipe(int *fildes*[2]);27814 **DESCRIPTION**

27815 The *pipe()* function shall create a pipe and place two file descriptors, one each into the  
 27816 arguments *fildes*[0] and *fildes*[1], that refer to the open file descriptions for the read and write  
 27817 ends of the pipe. Their integer values shall be the two lowest available at the time of the *pipe()*  
 27818 call. The O\_NONBLOCK and FD\_CLOEXEC flags shall be clear on both file descriptors. (The  
 27819 *fcntl()* function can be used to set both these flags.)

27820 Data can be written to the file descriptor *fildes*[1] and read from the file descriptor *fildes*[0]. A  
 27821 read on the file descriptor *fildes*[0] shall access data written to the file descriptor *fildes*[1] on a  
 27822 first-in-first-out basis. It is unspecified whether *fildes*[0] is also open for writing and whether  
 27823 *fildes*[1] is also open for reading.

27824 A process has the pipe open for reading (correspondingly writing) if it has a file descriptor open  
 27825 that refers to the read end, *fildes*[0] (write end, *fildes*[1]).

27826 Upon successful completion, *pipe()* shall mark for update the *st\_atime*, *st\_ctime*, and *st\_mtime*  
 27827 fields of the pipe.

27828 **RETURN VALUE**

27829 Upon successful completion, 0 shall be returned; otherwise, -1 shall be returned and *errno* set to  
 27830 indicate the error.

27831 **ERRORS**27832 The *pipe()* function shall fail if:

27833 [EMFILE] More than {OPEN\_MAX} minus two file descriptors are already in use by this  
 27834 process.

27835 [ENFILE] The number of simultaneously open files in the system would exceed a  
 27836 system-imposed limit.

27837 **EXAMPLES**

27838 None.

27839 **APPLICATION USAGE**

27840 None.

27841 **RATIONALE**

27842 The wording carefully avoids using the verb “to open” in order to avoid any implication of use  
 27843 of *open()*; see also *write()*.

27844 **FUTURE DIRECTIONS**

27845 None.

27846 **SEE ALSO**

27847 *fcntl()*, *read()*, *write()*, the Base Definitions volume of IEEE Std 1003.1-200x, <*fcntl.h*>,  
 27848 <*unistd.h*>

27849 **CHANGE HISTORY**

27850 First released in Issue 1. Derived from Issue 1 of the SVID.

27851 **Issue 6**

27852 The following new requirements on POSIX implementations derive from alignment with the  
27853 Single UNIX Specification:

- 27854 • The DESCRIPTION is updated to indicate that certain dispositions of *fildev[0]* and *fildev[1]*  
27855 are unspecified.

## 27856 NAME

27857 poll — input/output multiplexing

## 27858 SYNOPSIS

27859 XSI #include &lt;poll.h&gt;

27860 int poll(struct pollfd *fds*[], nfd\_t *nfds*, int *timeout*);

27861

## 27862 DESCRIPTION

27863 The *poll()* function provides applications with a mechanism for multiplexing input/output over  
 27864 a set of file descriptors. For each member of the array pointed to by *fds*, *poll()* shall examine the  
 27865 given file descriptor for the event(s) specified in *events*. The number of **pollfd** structures in the  
 27866 *fds* array is specified by *nfds*. The *poll()* function shall identify those file descriptors on which an  
 27867 application can read or write data, or on which certain events have occurred.

27868 The *fds* argument specifies the file descriptors to be examined and the events of interest for each  
 27869 file descriptor. It is a pointer to an array with one member for each open file descriptor of  
 27870 interest. The array's members are **pollfd** structures within which *fd* specifies an open file  
 27871 descriptor and *events* and *revents* are bitmasks constructed by OR'ing a combination of the  
 27872 following event flags:

27873	POLLIN	Data other than high-priority data may be read without blocking.
27874 XSR		For STREAMS, this flag is set in <i>revents</i> even if the message is of zero length.
27875		This flag shall be equivalent to POLLRDNORM   POLLRDBAND.
27876	POLLRDNORM	Normal data may be read without blocking.
27877 XSR		For STREAMS, data on priority band 0 may be read without blocking. This
27878		flag is set in <i>revents</i> even if the message is of zero length.
27879	POLLRDBAND	Priority data may be read without blocking.
27880 XSR		For STREAMS, data on priority bands greater than 0 may be read without
27881		blocking. This flag is set in <i>revents</i> even if the message is of zero length.
27882	POLLPRI	High-priority data may be read without blocking.
27883 XSR		For STREAMS, this flag is set in <i>revents</i> even if the message is of zero length.
27884	POLLOUT	Normal data may be written without blocking.
27885 XSR		For STREAMS, data on priority band 0 may be written without blocking.
27886	POLLWRNORM	Equivalent to POLLOUT.
27887	POLLWRBAND	Priority data may be written.
27888 XSR		For STREAMS, data on priority bands greater than 0 may be written without
27889		blocking. If any priority band has been written to on this STREAM, this event
27890		only examines bands that have been written to at least once.
27891	POLLERR	An error has occurred on the device or stream. This flag is only valid in the
27892		<i>revents</i> bitmask; it shall be ignored in the <i>events</i> member.
27893	POLLHUP	The device has been disconnected. This event and POLLOUT are mutually-
27894		exclusive; a stream can never be writable if a hangup has occurred. However,
27895		this event and POLLIN, POLLRDNORM, POLLRDBAND, or POLLPRI are not
27896		mutually-exclusive. This flag is only valid in the <i>revents</i> bitmask; it shall be
27897		ignored in the <i>events</i> member.

27898	POLLNVAL	The specified <i>fd</i> value is invalid. This flag is only valid in the <i>revents</i> member; it shall be ignored in the <i>events</i> member.
27899		
27900		The significance and semantics of normal, priority, and high-priority data are file and device-specific.
27901		
27902		If the value of <i>fd</i> is less than 0, <i>events</i> shall be ignored, and <i>revents</i> shall be set to 0 in that entry on return from <i>poll()</i> .
27903		
27904		In each <b>pollfd</b> structure, <i>poll()</i> shall clear the <i>revents</i> member, except that where the application requested a report on a condition by setting one of the bits of <i>events</i> listed above, <i>poll()</i> shall set the corresponding bit in <i>revents</i> if the requested condition is true. In addition, <i>poll()</i> shall set the POLLHUP, POLLERR, and POLLNVAL flag in <i>revents</i> if the condition is true, even if the application did not set the corresponding bit in <i>events</i> .
27905		
27906		
27907		
27908		
27909		If none of the defined events have occurred on any selected file descriptor, <i>poll()</i> shall wait at least <i>timeout</i> milliseconds for an event to occur on any of the selected file descriptors. If the value of <i>timeout</i> is 0, <i>poll()</i> shall return immediately. If the value of <i>timeout</i> is -1, <i>poll()</i> shall block until a requested event occurs or until the call is interrupted.
27910		
27911		
27912		
27913		Implementations may place limitations on the granularity of timeout intervals. If the requested timeout interval requires a finer granularity than the implementation supports, the actual timeout interval shall be rounded up to the next supported value.
27914		
27915		
27916		The <i>poll()</i> function shall not be affected by the O_NONBLOCK flag.
27917	XSR	The <i>poll()</i> function shall support regular files, terminal and pseudo-terminal devices, STREAMS-based files, FIFOs, pipes, and sockets. The behavior of <i>poll()</i> on elements of <i>fds</i> that refer to other types of file is unspecified.
27918		
27919		
27920		Regular files shall always poll TRUE for reading and writing.
27921		A file descriptor for a socket that is listening for connections shall indicate that it is ready for reading, once connections are available. A file descriptor for a socket that is connecting asynchronously shall indicate that it is ready for writing, once a connection has been established.
27922		
27923		
27924		<b>RETURN VALUE</b>
27925		Upon successful completion, <i>poll()</i> shall return a non-negative value. A positive value indicates the total number of file descriptors that have been selected (that is, file descriptors for which the <i>revents</i> member is non-zero). A value of 0 indicates that the call timed out and no file descriptors have been selected. Upon failure, <i>poll()</i> shall return -1 and set <i>errno</i> to indicate the error.
27926		
27927		
27928		
27929		<b>ERRORS</b>
27930		The <i>poll()</i> function shall fail if:
27931	[EAGAIN]	The allocation of internal data structures failed but a subsequent request may succeed.
27932		
27933	[EINTR]	A signal was caught during <i>poll()</i> .
27934	XSR	[EINVAL] The <i>nfds</i> argument is greater than {OPEN_MAX}, or one of the <i>fd</i> members refers to a STREAM or multiplexer that is linked (directly or indirectly) downstream from a multiplexer.
27935		
27936		



27937 **EXAMPLES**27938 **Checking for Events on a Stream**

27939 The following example opens a pair of STREAMS devices and then waits for either one to  
27940 become writable. This example proceeds as follows:

- 27941 1. Sets the *timeout* parameter to 500 milliseconds.
- 27942 2. Opens the STREAMS devices */dev/dev0* and */dev/dev1*, and then polls them, specifying  
27943 POLLOUT and POLLWRBAND as the events of interest.
- 27944 The STREAMS device names */dev/dev0* and */dev/dev1* are only examples of how  
27945 STREAMS devices can be named; STREAMS naming conventions may vary among  
27946 systems conforming to the IEEE Std 1003.1-200x.
- 27947 3. Uses the *ret* variable to determine whether an event has occurred on either of the two  
27948 STREAMS. The *poll()* function is given 500 milliseconds to wait for an event to occur (if it  
27949 has not occurred prior to the *poll()* call).
- 27950 4. Checks the returned value of *ret*. If a positive value is returned, one of the following can  
27951 be done:
  - 27952 a. Priority data can be written to the open STREAM on priority bands greater than 0,  
27953 because the POLLWRBAND event occurred on the open STREAM (*fds[0]* or *fds[1]*).
  - 27954 b. Data can be written to the open STREAM on priority-band 0, because the POLLOUT  
27955 event occurred on the open STREAM (*fds[0]* or *fds[1]*).
- 27956 5. If the returned value is not a positive value, permission to write data to the open STREAM  
27957 (on any priority band) is denied.
- 27958 6. If the POLLHUP event occurs on the open STREAM (*fds[0]* or *fds[1]*), the device on the  
27959 open STREAM has disconnected.

```

27960 #include <stropts.h>
27961 #include <poll.h>
27962 ...
27963 struct pollfd fds[2];
27964 int timeout_msecs = 500;
27965 int ret;
27966     int i;

27967 /* Open STREAMS device. */
27968 fds[0].fd = open("/dev/dev0", ...);
27969 fds[1].fd = open("/dev/dev1", ...);
27970     fds[0].events = POLLOUT | POLLWRBAND;
27971     fds[1].events = POLLOUT | POLLWRBAND;

27972 ret = poll(fds, 2, timeout_msecs);

27973 if (ret > 0) {
27974     /* An event on one of the fds has occurred. */
27975     for (i=0; i<2; i++) {
27976         if (fds[i].revents & POLLWRBAND) {
27977             /* Priority data may be written on device number i. */
27978             ...
27979         }
27980         if (fds[i].revents & POLLOUT) {

```

```

27981          /* Data may be written on device number i. */
27982      ...
27983          }
27984      if (fds[i].revents & POLLHUP) {
27985          /* A hangup has occurred on device number i. */
27986      ...
27987          }
27988      }
27989  }

```

27990 **APPLICATION USAGE**

27991 None.

27992 **RATIONALE**

27993 None.

27994 **FUTURE DIRECTIONS**

27995 None.

27996 **SEE ALSO**

27997 *getmsg()*, *putmsg()*, *read()*, *select()*, *write()*, the Base Definitions volume of IEEE Std 1003.1-200x,  
27998 <poll.h>, <stropts.h>, Section 2.6 (on page 488)

27999 **CHANGE HISTORY**

28000 First released in Issue 4, Version 2.

28001 **Issue 5**

28002 Moved from X/OPEN UNIX extension to BASE.

28003 The description of POLLWRBAND is updated.

28004 **Issue 6**

28005 Text referring to sockets is added to the DESCRIPTION.

28006 Text relating to the XSI STREAMS Option Group is marked. |

28007 The Open Group Corrigendum Unnn/nn is applied, updating the DESCRIPTION of |  
28008 POLLWRBAND. |

28009 **NAME**

28010 popen — initiate pipe streams to or from a process

28011 **SYNOPSIS**28012 cx `#include <stdio.h>`28013 `FILE *popen(const char *command, const char *mode);`

28014

28015 **DESCRIPTION**

28016 The *popen()* function shall execute the command specified by the string *command*. It shall create a  
 28017 pipe between the calling program and the executed command, and shall return a pointer to a  
 28018 stream that can be used to either read from or write to the pipe.

28019 The environment of the executed command shall be as if a child process were created within the  
 28020 *popen()* call using the *fork()* function, and the child invoked the *sh* utility using the call:

28021 `execl(shell_path, "sh", "-c", command, (char *)0);`28022 where *shell\_path* is an unspecified pathname for the *sh* utility.

28023 The *popen()* function shall ensure that any streams from previous *popen()* calls that remain open  
 28024 in the parent process are closed in the new child process.

28025 The *mode* argument to *popen()* is a string that specifies I/O mode:

- 28026 1. If *mode* is *r*, when the child process is started, its file descriptor `STDOUT_FILENO` shall be  
 28027 the writable end of the pipe, and the file descriptor *fileno(stream)* in the calling process,  
 28028 where *stream* is the stream pointer returned by *popen()*, shall be the readable end of the  
 28029 pipe.
- 28030 2. If *mode* is *w*, when the child process is started its file descriptor `STDIN_FILENO` shall be  
 28031 the readable end of the pipe, and the file descriptor *fileno(stream)* in the calling process,  
 28032 where *stream* is the stream pointer returned by *popen()*, shall be the writable end of the  
 28033 pipe.
- 28034 3. If *mode* is any other value, the result is undefined.

28035 After *popen()*, both the parent and the child process shall be capable of executing independently  
 28036 before either terminates.

28037 Pipe streams are byte-oriented.

28038 **RETURN VALUE**

28039 Upon successful completion, *popen()* shall return a pointer to an open stream that can be used to  
 28040 read or write to the pipe. Otherwise, it shall return a null pointer and may set *errno* to indicate  
 28041 the error.

28042 **ERRORS**28043 The *popen()* function may fail if:

28044 [EMFILE] {FOPEN\_MAX} or {STREAM\_MAX} streams are currently open in the calling  
 28045 process.

28046 [EINVAL] The *mode* argument is invalid.28047 The *popen()* function may also set *errno* values as described by *fork()* or *pipe()*.

28048 **EXAMPLES**

28049 None.

28050 **APPLICATION USAGE**

28051 Since open files are shared, a mode *r* command can be used as an input filter and a mode *w* |  
 28052 command as an output filter.

28053 Buffered reading before opening an input filter may leave the standard input of that filter  
 28054 mispositioned. Similar problems with an output filter may be prevented by careful buffer  
 28055 flushing; for example, with *flush()*.

28056 A stream opened by *popen()* should be closed by *pclose()*.

28057 The behavior of *popen()* is specified for values of *mode* of *r* and *w*. Other modes such as *rb* and  
 28058 *wb* might be supported by specific implementations, but these would not be portable features.  
 28059 Note that historical implementations of *popen()* only check to see if the first character of *mode* is  
 28060 *r*. Thus, a *mode* of *robert the robot* would be treated as *mode r*, and a *mode* of *anything else* would be  
 28061 treated as *mode w*.

28062 If the application calls *waitpid()* or *waitid()* with a *pid* argument greater than 0, and it still has a  
 28063 stream that was called with *popen()* open, it must ensure that *pid* does not refer to the process  
 28064 started by *popen()*.

28065 To determine whether or not the environment specified in the Shell and Utilities volume of  
 28066 IEEE Std 1003.1-200x is present, use the function call:

```
28067 sysconf(_SC_2_VERSION)
```

28068 (See *sysconf()*).

28069 **RATIONALE**

28070 The *popen()* function should not be used by programs that have set user (or group) ID privileges.  
 28071 The *fork()* and *exec* family of functions (except *execlp()* and *execvp()*), should be used instead.  
 28072 This prevents any unforeseen manipulation of the environment of the user that could cause  
 28073 execution of commands not anticipated by the calling program.

28074 If the original and *popen()*ed processes both intend to read or write or read and write a common  
 28075 file, and either will be using FILE-type C functions (*fread()*, *fwrite()*, and so on), the rules for  
 28076 sharing file handles must be observed (see Section 2.5.1 (on page 485)).

28077 Since open files are shared, a mode *r* argument can be used as an input filter and a mode *w* |  
 28078 argument as an output filter.

28079 The behavior of *popen()* is specified for modes of *r* and *w*. Other modes such as *rb* and *wb* might  
 28080 be supported by specific implementations, but these would not be portable features. Note that  
 28081 historical implementations of *popen()* only check to see if the first character of *mode* is '*r*'.  
 28082 Thus, a *mode* of *robert the robot* would be treated as *mode r*, and a *mode* of *anything else* would be  
 28083 treated as *mode w*.

28084 If the application calls *waitpid()* with a *pid* argument greater than zero, and it still has a  
 28085 *popen()*ed stream open, it must ensure that *pid* does not refer to the process started by *popen()*.

28086 **FUTURE DIRECTIONS**

28087 None.

28088 **SEE ALSO**

28089 *pclose()*, *pipe()*, *sysconf()*, *system()*, the Base Definitions volume of IEEE Std 1003.1-200x,  
 28090 **<stdio.h>**, the Shell and Utilities volume of IEEE Std 1003.1-200x, *sh* |

28091 **CHANGE HISTORY**

28092 First released in Issue 1. Derived from Issue 1 of the SVID.

28093 **Issue 5**

28094 A statement is added to the DESCRIPTION indicating that pipe streams are byte-oriented.

28095 **Issue 6**

28096 The following new requirements on POSIX implementations derive from alignment with the  
28097 Single UNIX Specification:

- 28098 • The optional [EMFILE] error condition is added.

28099 **NAME**

28100 posix\_fadvise — file advisory information (**ADVANCED REALTIME**)

28101 **SYNOPSIS**

28102 ADV #include <fcntl.h>

28103 int posix\_fadvise(int fd, off\_t offset, size\_t len, int advice);

28104

28105 **DESCRIPTION**

28106 The *posix\_fadvise()* function shall advise the implementation on the expected behavior of the  
 28107 application with respect to the data in the file associated with the open file descriptor, *fd*,  
 28108 starting at *offset* and continuing for *len* bytes. The specified range need not currently exist in the  
 28109 file. If *len* is zero, all data following *offset* is specified. The implementation may use this  
 28110 information to optimize handling of the specified data. The *posix\_fadvise()* function shall have no  
 28111 effect on the semantics of other operations on the specified data, although it may affect the  
 28112 performance of other operations.

28113 The advice to be applied to the data is specified by the *advice* parameter and may be one of the  
 28114 following values:

28115 POSIX\_FADV\_NORMAL

28116 Specifies that the application has no advice to give on its behavior with respect to the  
 28117 specified data. It is the default characteristic if no advice is given for an open file.

28118 POSIX\_FADV\_SEQUENTIAL

28119 Specifies that the application expects to access the specified data sequentially from lower  
 28120 offsets to higher offsets.

28121 POSIX\_FADV\_RANDOM

28122 Specifies that the application expects to access the specified data in a random order.

28123 POSIX\_FADV\_WILLNEED

28124 Specifies that the application expects to access the specified data in the near future.

28125 POSIX\_FADV\_DONTNEED

28126 Specifies that the application expects that it will not access the specified data in the near  
 28127 future.

28128 POSIX\_FADV\_NOREUSE

28129 Specifies that the application expects to access the specified data once and then not reuse it  
 28130 thereafter.

28131 These values are defined in <fcntl.h>.

28132 **RETURN VALUE**

28133 Upon successful completion, *posix\_fadvise()* shall return zero; otherwise, an error number shall  
 28134 be returned to indicate the error.

28135 **ERRORS**

28136 The *posix\_fadvise()* function shall fail if:

28137 [EBADF] The *fd* argument is not a valid file descriptor.

28138 [EINVAL] The value of *advice* is invalid.

28139 [ESPIPE] The *fd* argument is associated with a pipe or FIFO.

28140 **EXAMPLES**

28141           None.

28142 **APPLICATION USAGE**

28143           The *posix\_fadvise()* function is part of the Advisory Information option and need not be provided  
28144           on all implementations.

28145 **RATIONALE**

28146           None.

28147 **FUTURE DIRECTIONS**

28148           None.

28149 **SEE ALSO**28150           *posix\_madvise()*, the Base Definitions volume of IEEE Std 1003.1-200x, <fcntl.h>28151 **CHANGE HISTORY**

28152           First released in Issue 6. Derived from IEEE Std 1003.1d-1999.

28153           In the SYNOPSIS, the inclusion of &lt;sys/types.h&gt; is no longer required.

28154 **NAME**

28155 posix\_fallocate — file space control (**ADVANCED REALTIME**)

28156 **SYNOPSIS**

28157 ADV #include <fcntl.h>

28158 int posix\_fallocate(int fd, off\_t offset, size\_t len);

28159

28160 **DESCRIPTION**

28161 The *posix\_fallocate()* function shall ensure that any required storage for regular file data starting  
 28162 at *offset* and continuing for *len* bytes is allocated on the file system storage media. If  
 28163 *posix\_fallocate()* returns successfully, subsequent writes to the specified file data shall not fail  
 28164 due to the lack of free space on the file system storage media.

28165 If the *offset+len* is beyond the current file size, then *posix\_fallocate()* shall adjust the file size to  
 28166 *offset+len*. Otherwise, the file size shall not be changed.

28167 It is implementation-defined whether a previous *posix\_fadvise()* call influences allocation  
 28168 strategy.

28169 Space allocated via *posix\_fallocate()* shall be freed by a successful call to *creat()* or *open()* that  
 28170 truncates the size of the file. Space allocated via *posix\_fallocate()* may be freed by a successful call  
 28171 to *ftruncate()* that reduces the file size to a size smaller than *offset+len*.

28172 **RETURN VALUE**

28173 Upon successful completion, *posix\_fallocate()* shall return zero; otherwise, an error number shall  
 28174 be returned to indicate the error.

28175 **ERRORS**

28176 The *posix\_fallocate()* function shall fail if:

- 28177 [EBADF] The *fd* argument is not a valid file descriptor.
- 28178 [EBADF] The *fd* argument references a file that was opened without write permission.
- 28179 [EFBIG] The value of *offset+len* is greater than the maximum file size.
- 28180 [EINTR] A signal was caught during execution.
- 28181 [EINVAL] The *len* argument was zero or the *offset* argument was less than zero.
- 28182 [EIO] An I/O error occurred while reading from or writing to a file system.
- 28183 [ENODEV] The *fd* argument does not refer to a regular file.
- 28184 [ENOSPC] There is insufficient free space remaining on the file system storage media.
- 28185 [ESPIPE] The *fd* argument is associated with a pipe or FIFO.

28186 **EXAMPLES**

28187 None.

28188 **APPLICATION USAGE**

28189 The *posix\_fallocate()* function is part of the Advisory Information option and need not be  
 28190 provided on all implementations.

28191 **RATIONALE**

28192 None.



28193 **FUTURE DIRECTIONS**

28194       None.

28195 **SEE ALSO**28196       *creat()*, *truncate()*, *open()*, *unlink()*, the Base Definitions volume of IEEE Std 1003.1-200x,  
28197       <fcntl.h>28198 **CHANGE HISTORY**

28199       First released in Issue 6. Derived from IEEE Std 1003.1d-1999.

28200       In the SYNOPSIS, the inclusion of &lt;sys/types.h&gt; is no longer required.

28201 **NAME**

28202 posix\_madvise — memory advisory information and alignment control (**ADVANCED**  
28203 **REALTIME**)

28204 **SYNOPSIS**

28205 ADV #include <sys/mman.h>

28206 int posix\_madvise(void \*addr, size\_t len, int advice);  
28207

28208 **DESCRIPTION**

28209 MF|SHM The *posix\_madvise()* function need only be supported if either the Memory Mapped Files or the  
28210 Shared Memory Objects options are supported.

28211 The *posix\_madvise()* function shall advise the implementation on the expected behavior of the  
28212 application with respect to the data in the memory starting at address *addr*, and continuing for  
28213 *len* bytes. The implementation may use this information to optimize handling of the specified  
28214 data. The *posix\_madvise()* function shall have no effect on the semantics of access to memory in  
28215 the specified range, although it may affect the performance of access.

28216 The implementation may require that *addr* be a multiple of the page size, which is the value  
28217 returned by *sysconf()* when the name value *\_SC\_PAGESIZE* is used.

28218 The advice to be applied to the memory range is specified by the *advice* parameter and may be  
28219 one of the following values:

28220 POSIX\_MADV\_NORMAL

28221 Specifies that the application has no advice to give on its behavior with respect to the  
28222 specified range. It is the default characteristic if no advice is given for a range of memory.

28223 POSIX\_MADV\_SEQUENTIAL

28224 Specifies that the application expects to access the specified range sequentially from lower  
28225 addresses to higher addresses.

28226 POSIX\_MADV\_RANDOM

28227 Specifies that the application expects to access the specified range in a random order.

28228 POSIX\_MADV\_WILLNEED

28229 Specifies that the application expects to access the specified range in the near future.

28230 POSIX\_MADV\_DONTNEED

28231 Specifies that the application expects that it will not access the specified range in the near  
28232 future.

28233 These values are defined in the <sys/mman.h> header.

28234 **RETURN VALUE**

28235 Upon successful completion, *posix\_madvise()* shall return zero; otherwise, an error number shall  
28236 be returned to indicate the error.

28237 **ERRORS**

28238 The *posix\_madvise()* function shall fail if:

28239 [EINVAL] The value of *advice* is invalid.

28240 [ENOMEM] Addresses in the range starting at *addr* and continuing for *len* bytes are partly  
28241 or completely outside the range allowed for the address space of the calling  
28242 process.

28243 The *posix\_madvise()* function may fail if:

- 28244 [EINVAL] The value of *addr* is not a multiple of the value returned by *sysconf()* when the  
28245 name value `_SC_PAGESIZE` is used.
- 28246 [EINVAL] The value of *len* is zero.
- 28247 **EXAMPLES**
- 28248 None.
- 28249 **APPLICATION USAGE**
- 28250 The *posix\_madvise()* function is part of the Advisory Information option and need not be  
28251 provided on all implementations.
- 28252 **RATIONALE**
- 28253 None.
- 28254 **FUTURE DIRECTIONS**
- 28255 None.
- 28256 **SEE ALSO**
- 28257 *mmap()*, *posix\_fadvise()*, *sysconf()*, the Base Definitions volume of IEEE Std 1003.1-200x, |  
28258 `<sys/mman.h>` |
- 28259 **CHANGE HISTORY**
- 28260 First released in Issue 6. Derived from IEEE Std 1003.1d-1999.
- 28261 IEEE PASC Interpretation 1003.1 #102 is applied. |

28262 **NAME**

28263 posix\_mem\_offset — find offset and length of a mapped typed memory block (**ADVANCED**  
28264 **REALTIME**)

28265 **SYNOPSIS**

28266 TYM #include <sys/mman.h>

```
28267 int posix_mem_offset(const void *restrict addr, size_t len,
28268 off_t *restrict off, size_t *restrict contig_len,
28269 int *restrict fildes);
28270
```

28271 **DESCRIPTION**

28272 The *posix\_mem\_offset()* function shall return in the variable pointed to by *off* a value that  
28273 identifies the offset (or location), within a memory object, of the memory block currently  
28274 mapped at *addr*. The function shall return in the variable pointed to by *fildes*, the descriptor used  
28275 (via *mmap()*) to establish the mapping which contains *addr*. If that descriptor was closed since  
28276 the mapping was established, the returned value of *fildes* shall be  $-1$ . The *len* argument specifies  
28277 the length of the block of the memory object the user wishes the offset for; upon return, the value  
28278 pointed to by *contig\_len* shall equal either *len*, or the length of the largest contiguous block of the  
28279 memory object that is currently mapped to the calling process starting at *addr*, whichever is  
28280 smaller.

28281 If the memory object mapped at *addr* is a typed memory object, then if the *off* and *contig\_len*  
28282 values obtained by calling *posix\_mem\_offset()* are used in a call to *mmap()* with a file descriptor  
28283 that refers to the same memory pool as *fildes* (either through the same port or through a different  
28284 port), and that was opened with neither the `POSIX_TYPED_MEM_ALLOCATE` nor the  
28285 `POSIX_TYPED_MEM_ALLOCATE_CONTIG` flag, the typed memory area that is mapped shall  
28286 be exactly the same area that was mapped at *addr* in the address space of the process that called  
28287 *posix\_mem\_offset()*.

28288 If the memory object specified by *fildes* is not a typed memory object, then the behavior of this  
28289 function is implementation-defined.

28290 **RETURN VALUE**

28291 Upon successful completion, the *posix\_mem\_offset()* function shall return zero; otherwise, the  
28292 corresponding error status value shall be returned.

28293 **ERRORS**

28294 The *posix\_mem\_offset()* function shall fail if:

28295 [EACCES] The process has not mapped a memory object supported by this function at  
28296 the given address *addr*.

28297 This function shall not return an error code of [EINTR].

28298 **EXAMPLES**

28299 None.

28300 **APPLICATION USAGE**

28301 None.

28302 **RATIONALE**

28303 None.

28304 **FUTURE DIRECTIONS**

28305           None.

28306 **SEE ALSO**

28307           *mmap()*, *posix\_typed\_mem\_open()*, the Base Definitions volume of IEEE Std 1003.1-200x,  
28308           <**sys/mman.h**>

28309 **CHANGE HISTORY**

28310           First released in Issue 6. Derived from IEEE Std 1003.1j-2000.

28311 **NAME**

28312 posix\_memalign — aligned memory allocation (**ADVANCED REALTIME**)

28313 **SYNOPSIS**

28314 ADV `#include <stdlib.h>`

28315 `int posix_memalign(void **memptr, size_t alignment, size_t size);`

28316

28317 **DESCRIPTION**

28318 The *posix\_memalign()* function shall allocate *size* bytes aligned on a boundary specified by  
 28319 *alignment*, and shall return a pointer to the allocated memory in *memptr*. The value of *alignment*  
 28320 shall be a multiple of *sizeof(void \*)*, that is also a power of two. Upon successful completion, the  
 28321 value pointed to by *memptr* shall be a multiple of *alignment*.

28322 CX The *free()* function shall deallocate memory that has previously been allocated by  
 28323 *posix\_memalign()*.

28324 **RETURN VALUE**

28325 Upon successful completion, *posix\_memalign()* shall return zero; otherwise, an error number  
 28326 shall be returned to indicate the error.

28327 **ERRORS**

28328 The *posix\_memalign()* function shall fail if:

28329 [EINVAL] The value of the alignment parameter is not a power of two multiple of  
 28330 *sizeof(void \*)*.

28331 [ENOMEM] There is insufficient memory available with the requested alignment.

28332 **EXAMPLES**

28333 None.

28334 **APPLICATION USAGE**

28335 The *posix\_memalign()* function is part of the Advisory Information option and need not be  
 28336 provided on all implementations.

28337 **RATIONALE**

28338 None.

28339 **FUTURE DIRECTIONS**

28340 None.

28341 **SEE ALSO**

28342 *free()*, *malloc()*, the Base Definitions volume of IEEE Std 1003.1-200x, `<stdlib.h>`

28343 **CHANGE HISTORY**

28344 First released in Issue 6. Derived from IEEE Std 1003.1d-1999.

28345 In the SYNOPSIS, the inclusion of `<sys/types.h>` is no longer required.

28346 **NAME**

28347        posix\_openpt — open a pseudo terminal device

28348 **SYNOPSIS**

28349 XSI        #include &lt;stdlib.h&gt;

28350        #include &lt;fcntl.h&gt;

28351        int posix\_openpt(int oflag);

28352

28353 **DESCRIPTION**

28354        The *posix\_openpt()* function shall establish a connection between a master device for a pseudo-terminal and a file descriptor. The file descriptor is used by other I/O functions that refer to that pseudo-terminal.

28357        The file status flags and file access modes of the open file description shall be set according to the value of *oflag*.

28359        Values for *oflag* are constructed by a bitwise-inclusive OR of flags from the following list, defined in <fcntl.h>:

28361        O\_RDWR        Open for reading and writing.

28362        O\_NOCTTY       If set *posix\_openpt()* shall not cause the terminal device to become the controlling terminal for the process.

28364        The behavior of other values for the *oflag* argument is unspecified.

28365 **RETURN VALUE**

28366        Upon successful completion, the *posix\_openpt()* function shall open a master pseudo-terminal device and return a non-negative integer representing the lowest numbered unused file descriptor. Otherwise, -1 shall be returned and *errno* set to indicate the error.

28369 **ERRORS**

28370        The *posix\_openpt()* function shall fail if:

28371        [EMFILE]       {OPEN\_MAX} file descriptors are currently open in the calling process.

28372        [ENFILE]       The maximum allowable number of files is currently open in the system.

28373        The *posix\_openpt()* function may fail if:

28374        [EINVAL]       The value of *oflag* is not valid.

28375        [EAGAIN]       Out of pseudo-terminal resources.

28376 XSR        [ENOSR]       Out of STREAMS resources.

28377 **EXAMPLES**

28378        **Opening a Pseudo-Terminal and Returning the Name of the Slave Device and a File Descriptor**

28380        #include <fcntl.h>

28381        #include <stdio.h>

28382        int masterfd, slavefd;

28383        char \*slavedevice;

28384        masterfd = posix\_openpt(O\_RDWR|O\_NOCTTY);

28385        if (masterfd == -1

28386            || grantpt (masterfd) == -1

```

28387         || unlockpt (masterfd) == -1
28388         || (slavedevice = ptsname (masterfd)) == NULL)
28389         return -1;

28390     printf("slave device is: %s\n", slavedevice);

28391     slavefd = open(slave, O_RDWR|O_NOCTTY);
28392     if (slavefd < 0)
28393         return -1;

```

28394 **APPLICATION USAGE**

28395 This function is a method for portably obtaining a file descriptor of a master terminal device for a pseudo-terminal. The *grantpt()* and *ptsname()* functions can be used to manipulate mode and ownership permissions, and to obtain the name of the slave device, respectively.

28398 **RATIONALE**

28399 The standard developers considered the matter of adding a special device for cloning master pseudo-terminals: the */dev/ptmx* device. However, consensus could not be reached, and it was felt that adding a new function would permit other implementations. The *posix\_openpt()* function is designed to complement the *grantpt()*, *ptsname()*, and *unlockpt()* functions.

28403 On implementations supporting the */dev/ptmx* clone device, opening the master device of a pseudo-terminal is simply:

```

28405     mfdp = open("/dev/ptmx", oflag );
28406     if (mfdp < 0)
28407         return -1;

```

28408 **FUTURE DIRECTIONS**

28409 None.

28410 **SEE ALSO**

28411 *grantpt()*, *open()*, *ptsname()*, *unlockpt()*, the Base Definitions volume of IEEE Std 1003.1-200x, <fcntl.h>

28413 **CHANGE HISTORY**

28414 First released in Issue 6.



28415 **NAME**28416 posix\_spawn, posix\_spawnnp — spawn a process (**ADVANCED REALTIME**)28417 **SYNOPSIS**

```
28418 SPN #include <spawn.h>
28419
28419 int posix_spawn(pid_t *restrict pid, const char *restrict path,
28420               const posix_spawn_file_actions_t *file_actions,
28421               const posix_spawnattr_t *restrict attrp,
28422               char *const argv[restrict], char *const envp[restrict]);
28423 int posix_spawnnp(pid_t *restrict pid, const char *restrict file,
28424                 const posix_spawn_file_actions_t *file_actions,
28425                 const posix_spawnattr_t *restrict attrp,
28426                 char *const argv[restrict], char * const envp[restrict]);
28427
```

28428 **DESCRIPTION**

28429 The *posix\_spawn()* and *posix\_spawnnp()* functions shall create a new process (child process) from  
 28430 the specified process image. The new process image shall be constructed from a regular  
 28431 executable file called the new process image file.

28432 When a C program is executed as the result of this call, it shall be entered as a C language  
 28433 function call as follows:

```
28434 int main(int argc, char *argv[]);
```

28435 where *argc* is the argument count and *argv* is an array of character pointers to the arguments  
 28436 themselves. In addition, the following variable:

```
28437 extern char **environ;
```

28438 shall be initialized as a pointer to an array of character pointers to the environment strings.

28439 The argument *argv* is an array of character pointers to null-terminated strings. The last member  
 28440 of this array shall be a null pointer and is not counted in *argc*. These strings constitute the  
 28441 argument list available to the new process image. The value in *argv*[0] should point to a filename  
 28442 that is associated with the process image being started by the *posix\_spawn()* or *posix\_spawnnp()*  
 28443 function.

28444 The argument *envp* is an array of character pointers to null-terminated strings. These strings  
 28445 constitute the environment for the new process image. The environment array is terminated by a  
 28446 null pointer.

28447 The number of bytes available for the child process' combined argument and environment lists  
 28448 is {ARG\_MAX}. The implementation shall specify in the system documentation (see the Base  
 28449 Definitions volume of IEEE Std 1003.1-200x, Chapter 2, Conformance) whether any list  
 28450 overhead, such as length words, null terminators, pointers, or alignment bytes, is included in  
 28451 this total.

28452 The *path* argument to *posix\_spawn()* is a pathname that identifies the new process image file to  
 28453 execute.

28454 The *file* parameter to *posix\_spawnnp()* shall be used to construct a pathname that identifies the  
 28455 new process image file. If the *file* parameter contains a slash character, the *file* parameter shall be  
 28456 used as the pathname for the new process image file. Otherwise, the path prefix for this file shall  
 28457 be obtained by a search of the directories passed as the environment variable *PATH* (see the Base  
 28458 Definitions volume of IEEE Std 1003.1-200x, Chapter 8, Environment Variables). If this  
 28459 environment variable is not defined, the results of the search are implementation-defined.

28460 If *file\_actions* is a null pointer, then file descriptors open in the calling process shall remain open  
 28461 in the child process, except for those whose close-on-exec flag FD\_CLOEXEC is set (see *fcntl()*).  
 28462 For those file descriptors that remain open, all attributes of the corresponding open file  
 28463 descriptions, including file locks (see *fcntl()*), shall remain unchanged.

28464 If *file\_actions* is not NULL, then the file descriptors open in the child process shall be those open  
 28465 in the calling process as modified by the spawn file actions object pointed to by *file\_actions* and  
 28466 the FD\_CLOEXEC flag of each remaining open file descriptor after the spawn file actions have  
 28467 been processed. The effective order of processing the spawn file actions shall be:

- 28468 1. The set of open file descriptors for the child process shall initially be the same set as is  
 28469 open for the calling process. All attributes of the corresponding open file descriptions,  
 28470 including file locks (see *fcntl()*), shall remain unchanged.
- 28471 2. The signal mask, signal default actions, and the effective user and group IDs for the child  
 28472 process shall be changed as specified in the attributes object referenced by *attrp*.
- 28473 3. The file actions specified by the spawn file actions object shall be performed in the order in  
 28474 which they were added to the spawn file actions object.
- 28475 4. Any file descriptor that has its FD\_CLOEXEC flag set (see *fcntl()*) shall be closed.

28476 The **posix\_spawnattr\_t** spawn attributes object type is defined in **<spawn.h>**. It shall contain at  
 28477 least the attributes defined below.

28478 If the POSIX\_SPAWN\_SETPGROUP flag is set in the *spawn\_flags* attribute of the object  
 28479 referenced by *attrp*, and the *spawn-pgroup* attribute of the same object is non-zero, then the  
 28480 child's process group shall be as specified in the *spawn-pgroup* attribute of the object referenced  
 28481 by *attrp*.

28482 As a special case, if the POSIX\_SPAWN\_SETPGROUP flag is set in the *spawn\_flags* attribute of  
 28483 the object referenced by *attrp*, and the *spawn-pgroup* attribute of the same object is set to zero,  
 28484 then the child shall be in a new process group with a process group ID equal to its process ID.

28485 If the POSIX\_SPAWN\_SETPGROUP flag is not set in the *spawn\_flags* attribute of the object  
 28486 referenced by *attrp*, the new child process shall inherit the parent's process group.

28487 PS If the POSIX\_SPAWN\_SETSCHEDPARAM flag is set in the *spawn\_flags* attribute of the object  
 28488 referenced by *attrp*, but POSIX\_SPAWN\_SETSCHEDULER is not set, the new process image  
 28489 shall initially have the scheduling policy of the calling process with the scheduling parameters  
 28490 specified in the *spawn-schedparam* attribute of the object referenced by *attrp*.

28491 If the POSIX\_SPAWN\_SETSCHEDULER flag is set in *spawn\_flags* attribute of the object  
 28492 referenced by *attrp* (regardless of the setting of the POSIX\_SPAWN\_SETSCHEDPARAM flag),  
 28493 the new process image shall initially have the scheduling policy specified in the *spawn-*  
 28494 *schedpolicy* attribute of the object referenced by *attrp* and the scheduling parameters specified in  
 28495 the *spawn-schedparam* attribute of the same object.

28496 The POSIX\_SPAWN\_RESETPID flag in the *spawn\_flags* attribute of the object referenced by *attrp* |  
 28497 governs the effective user ID of the child process. If this flag is not set, the child process shall |  
 28498 inherit the parent process' effective user ID. If this flag is set, the child process' effective user ID |  
 28499 shall be reset to the parent's real user ID. In either case, if the set-user-ID mode bit of the new |  
 28500 process image file is set, the effective user ID of the child process shall become that file's owner |  
 28501 ID before the new process image begins execution. |

28502 The POSIX\_SPAWN\_RESETPID flag in the *spawn\_flags* attribute of the object referenced by *attrp* |  
 28503 also governs the effective group ID of the child process. If this flag is not set, the child process |  
 28504 shall inherit the parent process' effective group ID. If this flag is set, the child process' effective |  
 28505 group ID shall be reset to the parent's real group ID. In either case, if the set-group-ID mode bit |

28506 of the new process image file is set, the effective group ID of the child process shall become that  
28507 file's group ID before the new process image begins execution.

28508 If the POSIX\_SPAWN\_SETSIGMASK flag is set in the *spawn\_flags* attribute of the object  
28509 referenced by *attrp*, the child process shall initially have the signal mask specified in the *spawn-*  
28510 *sigmask* attribute of the object referenced by *attrp*.

28511 If the POSIX\_SPAWN\_SETSIGDEF flag is set in the *spawn\_flags* attribute of the object referenced  
28512 by *attrp*, the signals specified in the *spawn\_sigdefault* attribute of the same object shall be set to  
28513 their default actions in the child process. Signals set to the default action in the parent process  
28514 shall be set to the default action in the child process.

28515 Signals set to be caught by the calling process shall be set to the default action in the child  
28516 process.

28517 Except for SIGCHLD, signals set to be ignored by the calling process image shall be set to be |  
28518 ignored by the child process, unless otherwise specified by the POSIX\_SPAWN\_SETSIGDEF flag |  
28519 being set in the *spawn\_flags* attribute of the object referenced by *attrp* and the signals being |  
28520 indicated in the *spawn\_sigdefault* attribute of the object referenced by *attrp*.

28521 If the SIGCHLD signal is set to be ignored by the calling process, it is unspecified whether the |  
28522 SIGCHLD signal is set to be ignored or to the default action in the child process, unless |  
28523 otherwise specified by the POSIX\_SPAWN\_SETSIGDEF flag being set in the *spawn\_flags* |  
28524 attribute of the object referenced by *attrp* and the SIGCHLD signal being indicated in the |  
28525 *spawn\_sigdefault* attribute of the object referenced by *attrp*.

28526 If the value of the *attrp* pointer is NULL, then the default values are used.

28527 All process attributes, other than those influenced by the attributes set in the object referenced  
28528 by *attrp* as specified above or by the file descriptor manipulations specified in *file\_actions*, shall  
28529 appear in the new process image as though *fork()* had been called to create a child process and  
28530 then a member of the *exec* family of functions had been called by the child process to execute the  
28531 new process image.

28532 THR It is implementation-defined whether the fork handlers are run when *posix\_spawn()* or  
28533 *posix\_spawnnp()* is called.

#### 28534 RETURN VALUE

28535 Upon successful completion, *posix\_spawn()* and *posix\_spawnnp()* shall return the process ID of the  
28536 child process to the parent process, in the variable pointed to by a non-NULL *pid* argument, and  
28537 shall return zero as the function return value. Otherwise, no child process shall be created, the  
28538 value stored into the variable pointed to by a non-NULL *pid* is unspecified, and an error number  
28539 shall be returned as the function return value to indicate the error. If the *pid* argument is a null  
28540 pointer, the process ID of the child is not returned to the caller.

#### 28541 ERRORS

28542 The *posix\_spawn()* and *posix\_spawnnp()* functions may fail if:

28543 [EINVAL] The value specified by *file\_actions* or *attrp* is invalid.

28544 If this error occurs after the calling process successfully returns from the *posix\_spawn()* or  
28545 *posix\_spawnnp()* function, the child process may exit with exit status 127.

28546 If *posix\_spawn()* or *posix\_spawnnp()* fail for any of the reasons that would cause *fork()* or one of  
28547 the *exec* family of functions to fail, an error value shall be returned as described by *fork()* and  
28548 *exec*, respectively (or, if the error occurs after the calling process successfully returns, the child  
28549 process shall exit with exit status 127).

28550 If POSIX\_SPAWN\_SETPGROUP is set in the *spawn-flags* attribute of the object referenced by  
 28551 *attrp*, and *posix\_spawn()* or *posix\_spawnp()* fails while changing the child's process group, an  
 28552 error value shall be returned as described by *setpgid()* (or, if the error occurs after the calling  
 28553 process successfully returns, the child process shall exit with exit status 127).

28554 PS If POSIX\_SPAWN\_SETSCHEDPARAM is set and POSIX\_SPAWN\_SETSCHEDULER is not set  
 28555 in the *spawn-flags* attribute of the object referenced by *attrp*, then if *posix\_spawn()* or  
 28556 *posix\_spawnp()* fails for any of the reasons that would cause *sched\_setparam()* to fail, an error  
 28557 value shall be returned as described by *sched\_setparam()* (or, if the error occurs after the calling  
 28558 process successfully returns, the child process shall exit with exit status 127).

28559 If POSIX\_SPAWN\_SETSCHEDULER is set in the *spawn-flags* attribute of the object referenced by  
 28560 *attrp*, and if *posix\_spawn()* or *posix\_spawnp()* fails for any of the reasons that would cause  
 28561 *sched\_setscheduler()* to fail, an error value shall be returned as described by *sched\_setscheduler()*  
 28562 (or, if the error occurs after the calling process successfully returns, the child process shall exit  
 28563 with exit status 127).

28564 If the *file\_actions* argument is not NULL, and specifies any *close*, *dup2*, or *open* actions to be  
 28565 performed, and if *posix\_spawn()* or *posix\_spawnp()* fails for any of the reasons that would cause  
 28566 *close()*, *dup2()*, or *open()* to fail, an error value shall be returned as described by *close()*,  
 28567 *dup2()*, and *open()*, respectively (or, if the error occurs after the calling process successfully returns, the  
 28568 child process shall exit with exit status 127). An open file action may, by itself, result in any of  
 28569 the errors described by *close()* or *dup2()*, in addition to those described by *open()*.

28570 **EXAMPLES**

28571 None.

28572 **APPLICATION USAGE**

28573 These functions are part of the Spawn option and need not be provided on all implementations.

28574 **RATIONALE**

28575 The *posix\_spawn()* function and its close relation *posix\_spawnp()* have been introduced to |  
 28576 overcome the following perceived difficulties with *fork()*: the *fork()* function is difficult or |  
 28577 impossible to implement without swapping or dynamic address translation. |

- 28578 • Swapping is generally too slow for a realtime environment. |
- 28579 • Dynamic address translation is not available everywhere that POSIX might be useful. |
- 28580 • Processes are too useful to simply option out of POSIX whenever it must run without |  
 28581 address translation or other MMU services, |

28582 Thus, POSIX needs process creation and file execution primitives that can be efficiently |  
 28583 implemented without address translation or other MMU services. |

28584 The *posix\_spawn()* function is implementable as a library routine, but both *posix\_spawn()* and |  
 28585 *posix\_spawnp()* are designed as kernel operations. Also, although they may be an efficient |  
 28586 replacement for many *fork()/exec* pairs, their goal is to provide useful process creation |  
 28587 primitives for systems that have difficulty with *fork()*, not to provide drop-in replacements for |  
 28588 *fork()/exec*.

28589 This view of the role of *posix\_spawn()* and *posix\_spawnp()* influenced the design of their API. It  
 28590 does not attempt to provide the full functionality of *fork()/exec* in which arbitrary user-specified  
 28591 operations of any sort are permitted between the creation of the child process and the execution  
 28592 of the new process image; any attempt to reach that level would need to provide a programming  
 28593 language as parameters. Instead, *posix\_spawn()* and *posix\_spawnp()* are process creation  
 28594 primitives like the *Start\_Process* and *Start\_Process\_Search* Ada language bindings package  
 28595 *POSIX\_Process\_Primitives* and also like those in many operating systems that are not UNIX

28596 systems, but with some POSIX-specific additions.

28597 To achieve its coverage goals, *posix\_spawn()* and *posix\_spawnp()* have control of six types of  
 28598 inheritance: file descriptors, process group ID, user and group ID, signal mask, scheduling, and  
 28599 whether each signal ignored in the parent will remain ignored in the child, or be reset to its  
 28600 default action in the child.

28601 Control of file descriptors is required to allow an independently written child process image to  
 28602 access data streams opened by and even generated or read by the parent process without being  
 28603 specifically coded to know which parent files and file descriptors are to be used. Control of the  
 28604 process group ID is required to control how the child process' job control relates to that of the  
 28605 parent.

28606 Control of the signal mask and signal defaulting is sufficient to support the implementation of  
 28607 *system()*. Although support for *system()* is not explicitly one of the goals for *posix\_spawn()* and  
 28608 *posix\_spawnp()*, it is covered under the "at least 50%" coverage goal.

28609 The intention is that the normal file descriptor inheritance across *fork()*, the subsequent effect of  
 28610 the specified spawn file actions, and the normal file descriptor inheritance across one of the *exec*  
 28611 family of functions should fully specify open file inheritance. The implementation need make no  
 28612 decisions regarding the set of open file descriptors when the child process image begins  
 28613 execution, those decisions having already been made by the caller and expressed as the set of  
 28614 open file descriptors and their *FD\_CLOEXEC* flags at the time of the call and the spawn file  
 28615 actions object specified in the call. We have been assured that in cases where the POSIX  
 28616 *Start\_Process* Ada primitives have been implemented in a library, this method of controlling file  
 28617 descriptor inheritance may be implemented very easily.

28618 We can identify several problems with *posix\_spawn()* and *posix\_spawnp()*, but there does not  
 28619 appear to be a solution that introduces fewer problems. Environment modification for child  
 28620 process attributes not specifiable via the *attrp* or *file\_actions* arguments must be done in the  
 28621 parent process, and since the parent generally wants to save its context, it is more costly than  
 28622 similar functionality with *fork()/exec*. It is also complicated to modify the environment of a  
 28623 multi-threaded process temporarily, since all threads must agree when it is safe for the  
 28624 environment to be changed. However, this cost is only borne by those invocations of  
 28625 *posix\_spawn()* and *posix\_spawnp()* that use the additional functionality. Since extensive  
 28626 modifications are not the usual case, and are particularly unlikely in time-critical code, keeping  
 28627 much of the environment control out of *posix\_spawn()* and *posix\_spawnp()* is appropriate design.

28628 The *posix\_spawn()* and *posix\_spawnp()* functions do not have all the power of *fork()/exec*. This is  
 28629 to be expected. The *fork()* function is a wonderfully powerful operation. We do not expect to  
 28630 duplicate its functionality in a simple, fast function with no special hardware requirements. It is  
 28631 worth noting that *posix\_spawn()* and *posix\_spawnp()* are very similar to the process creation  
 28632 operations on many operating systems that are not UNIX systems.

### 28633 **Requirements**

28634 The requirements for *posix\_spawn()* and *posix\_spawnp()* are:

- 28635 • They must be implementable without an MMU or unusual hardware.
- 28636 • They must be compatible with existing POSIX standards.

28637 Additional goals are:

- 28638 • They should be efficiently implementable.
- 28639 • They should be able to replace at least 50% of typical executions of *fork()*.

28640           • A system with *posix\_spawn()* and *posix\_spawnp()* and without *fork()* should be useful, at least  
28641           for realtime applications.

28642           • A system with *fork()* and the *exec* family should be able to implement *posix\_spawn()* and  
28643           *posix\_spawnp()* as library routines.

28644           **Two-Syntax**

28645           POSIX *exec* has several calling sequences with approximately the same functionality. These  
28646           appear to be required for compatibility with existing practice. Since the existing practice for the  
28647           *posix\_spawn\*()* functions is otherwise substantially unlike POSIX, we feel that simplicity  
28648           outweighs compatibility. There are, therefore, only two names for the *posix\_spawn\*()* functions.

28649           The parameter list does not differ between *posix\_spawn()* and *posix\_spawnp()*; *posix\_spawnp()*  
28650           interprets the second parameter more elaborately than *posix\_spawn()*.

28651           **Compatibility with POSIX.5 (Ada)**

28652           The *Start\_Process* and *Start\_Process\_Search* procedures from the POSIX\_Process\_Primitives  
28653           package from the Ada language binding to POSIX.1 encapsulate *fork()* and *exec* functionality in a  
28654           manner similar to that of *posix\_spawn()* and *posix\_spawnp()*. Originally, in keeping with our  
28655           simplicity goal, the standard developers had limited the capabilities of *posix\_spawn()* and  
28656           *posix\_spawnp()* to a subset of the capabilities of *Start\_Process* and *Start\_Process\_Search*; certain  
28657           non-default capabilities were not supported. However, based on suggestions by the ballot group  
28658           to improve file descriptor mapping or drop it, and on the advice of an Ada Language Bindings  
28659           working group member, the standard developers decided that *posix\_spawn()* and *posix\_spawnp()*  
28660           should be sufficiently powerful to implement *Start\_Process* and *Start\_Process\_Search*. The  
28661           rationale is that if the Ada language binding to such a primitive had already been approved as  
28662           an IEEE standard, there can be little justification for not approving the functionally-equivalent  
28663           parts of a C binding. The only three capabilities provided by *posix\_spawn()* and *posix\_spawnp()*  
28664           that are not provided by *Start\_Process* and *Start\_Process\_Search* are optionally specifying the  
28665           child's process group ID, the set of signals to be reset to default signal handling in the child  
28666           process, and the child's scheduling policy and parameters.

28667           For the Ada language binding for *Start\_Process* to be implemented with *posix\_spawn()*, that  
28668           binding would need to explicitly pass an empty signal mask and the parent's environment to  
28669           *posix\_spawn()* whenever the caller of *Start\_Process* allowed these arguments to default, since  
28670           *posix\_spawn()* does not provide such defaults. The ability of *Start\_Process* to mask user-specified  
28671           signals during its execution is functionally unique to the Ada language binding and must be  
28672           dealt with in the binding separately from the call to *posix\_spawn()*.

28673           **Process Group**

28674           The process group inheritance field can be used to join the child process with an existing process  
28675           group. By assigning a value of zero to the *spawn-pgroup* attribute of the object referenced by  
28676           *attrp*, the *setpgid()* mechanism will place the child process in a new process group.

28677

**Threads**

28678

28679

28680

28681

28682

28683

28684

28685

28686

Without the *posix\_spawn()* and *posix\_spawnp()* functions, systems without address translation can still use threads to give an abstraction of concurrency. In many cases, thread creation suffices, but it is not always a good substitute. The *posix\_spawn()* and *posix\_spawnp()* functions are considerably “heavier” than thread creation. Processes have several important attributes that threads do not. Even without address translation, a process may have base-and-bound memory protection. Each process has a process environment including security attributes and file capabilities, and powerful scheduling attributes. Processes abstract the behavior of non-uniform-memory-architecture multi-processors better than threads, and they are more convenient to use for activities that are not closely linked.

28687

28688

28689

28690

28691

28692

The *posix\_spawn()* and *posix\_spawnp()* functions may not bring support for multiple processes to every configuration. Process creation is not the only piece of operating system support required to support multiple processes. The total cost of support for multiple processes may be quite high in some circumstances. Existing practice shows that support for multiple processes is uncommon and threads are common among “tiny kernels”. There should, therefore, probably continue to be AEPs for operating systems with only one process.

28693

**Asynchronous Error Notification**

28694

28695

28696

28697

28698

A library implementation of *posix\_spawn()* or *posix\_spawnp()* may not be able to detect all possible errors before it forks the child process. IEEE Std 1003.1-200x provides for an error indication returned from a child process which could not successfully complete the spawn operation via a special exit status which may be detected using the status value returned by *wait()* and *waitpid()*.

28699

28700

28701

28702

28703

The *stat\_val* interface and the macros used to interpret it are not well suited to the purpose of returning API errors, but they are the only path available to a library implementation. Thus, an implementation may cause the child process to exit with exit status 127 for any error detected during the spawn process after the *posix\_spawn()* or *posix\_spawnp()* function has successfully returned.

28704

28705

28706

28707

28708

28709

28710

28711

The standard developers had proposed using two additional macros to interpret *stat\_val*. The first, WIFSPAWNFAIL, would have detected a status that indicated that the child exited because of an error detected during the *posix\_spawn()* or *posix\_spawnp()* operations rather than during actual execution of the child process image; the second, WSPAWNERRNO, would have extracted the error value if WIFSPAWNFAIL indicated a failure. Unfortunately, the ballot group strongly opposed this because it would make a library implementation of *posix\_spawn()* or *posix\_spawnp()* dependent on kernel modifications to *waitpid()* to be able to embed special information in *stat\_val* to indicate a spawn failure.

28712

28713

28714

28715

28716

28717

28718

28719

28720

28721

28722

The 8 bits of child process exit status that are guaranteed by IEEE Std 1003.1-200x to be accessible to the waiting parent process are insufficient to disambiguate a spawn error from any other kind of error that may be returned by an arbitrary process image. No other bits of the exit status are required to be visible in *stat\_val*, so these macros could not be strictly implemented at the library level. Reserving an exit status of 127 for such spawn errors is consistent with the use of this value by *system()* and *popen()* to signal failures in these operations that occur after the function has returned but before a shell is able to execute. The exit status of 127 does not uniquely identify this class of error, nor does it provide any detailed information on the nature of the failure. Note that a kernel implementation of *posix\_spawn()* or *posix\_spawnp()* is permitted (and encouraged) to return any possible error as the function value, thus providing more detailed failure information to the parent process.

28723

28724

Thus, no special macros are available to isolate asynchronous *posix\_spawn()* or *posix\_spawnp()* errors. Instead, errors detected by the *posix\_spawn()* or *posix\_spawnp()* operations in the context

28725 of the child process before the new process image executes are reported by setting the child's  
 28726 exit status to 127. The calling process may use the WIFEXITED and WEXITSTATUS macros on  
 28727 the *stat\_val* stored by the *wait()* or *waitpid()* functions to detect spawn failures to the extent that  
 28728 other status values with which the child process image may exit (before the parent can  
 28729 conclusively determine that the child process image has begun execution) are distinct from exit  
 28730 status 127.

28731 **FUTURE DIRECTIONS**

28732 None.

28733 **SEE ALSO**

28734 *alarm()*, *chmod()*, *close()*, *dup()*, *exec*, *exit()*, *fcntl()*, *fork()*, *kill()*, *open()*,  
 28735 *posix\_spawn\_file\_actions\_addclose()*, *posix\_spawn\_file\_actions\_adddup2()*,  
 28736 *posix\_spawn\_file\_actions\_addopen()*, *posix\_spawn\_file\_actions\_destroy()*,  
 28737 *posix\_spawn\_file\_actions\_init()*, *posix\_spawnattr\_destroy()*, *posix\_spawnattr\_init()*,  
 28738 *posix\_spawnattr\_getsigdefault()*, *posix\_spawnattr\_getflags()*, *posix\_spawnattr\_getpgroup()*,  
 28739 *posix\_spawnattr\_getschedparam()*, *posix\_spawnattr\_getschedpolicy()*, *posix\_spawnattr\_getsigmask()*,  
 28740 *posix\_spawnattr\_setsigdefault()*, *posix\_spawnattr\_setflags()*, *posix\_spawnattr\_setpgroup()*,  
 28741 *posix\_spawnattr\_setschedparam()*, *posix\_spawnattr\_setschedpolicy()*, *posix\_spawnattr\_setsigmask()*,  
 28742 *sched\_setparam()*, *sched\_setscheduler()*, *setpgid()*, *setuid()*, *stat()*, *times()*, *wait()*, the Base  
 28743 Definitions volume of IEEE Std 1003.1-200x, <**spawn.h**>

28744 **CHANGE HISTORY**

28745 First released in Issue 6. Derived from IEEE Std 1003.1d-1999.

28746 IEEE PASC Interpretation 1003.1 #103 is applied, noting that the signal default actions are |  
 28747 changed as well as the signal mask in step 2. |

28748 IEEE PASC Interpretation 1003.1 #132 is applied. |



28749 **NAME**

28750 posix\_spawn\_file\_actions\_addclose, posix\_spawn\_file\_actions\_addopen — add close or open  
 28751 action to spawn file actions object (**ADVANCED REALTIME**)

28752 **SYNOPSIS**

```
28753 SPN #include <spawn.h>
28754 int posix_spawn_file_actions_addclose(posix_spawn_file_actions_t *
28755     file_actions, int fildes);
28756 int posix_spawn_file_actions_addopen(posix_spawn_file_actions_t *restrict
28757     file_actions, int fildes, const char *restrict path,
28758     int oflag, mode_t mode);
28759
```

28760 **DESCRIPTION**

28761 These functions shall add or delete a close or open action to a spawn file actions object. |

28762 A spawn file actions object is of type **posix\_spawn\_file\_actions\_t** (defined in `<spawn.h>`) and is |  
 28763 used to specify a series of actions to be performed by a `posix_spawn()` or `posix_spawnp()` |  
 28764 operation in order to arrive at the set of open file descriptors for the child process given the set of |  
 28765 open file descriptors of the parent. IEEE Std 1003.1-200x does not define comparison or |  
 28766 assignment operators for the type **posix\_spawn\_file\_actions\_t**.

28767 A spawn file actions object, when passed to `posix_spawn()` or `posix_spawnp()`, shall specify how |  
 28768 the set of open file descriptors in the calling process is transformed into a set of potentially open |  
 28769 file descriptors for the spawned process. This transformation shall be as if the specified sequence |  
 28770 of actions was performed exactly once, in the context of the spawned process (prior to execution |  
 28771 of the new process image), in the order in which the actions were added to the object; |  
 28772 additionally, when the new process image is executed, any file descriptor (from this new set) |  
 28773 which has its `FD_CLOEXEC` flag set shall be closed (see `posix_spawn()`). |

28774 The `posix_spawn_file_actions_addclose()` function shall add a *close* action to the object referenced |  
 28775 by `file_actions` that shall cause the file descriptor `fildes` to be closed (as if `close(fildes)` had been |  
 28776 called) when a new process is spawned using this file actions object.

28777 The `posix_spawn_file_actions_addopen()` function shall add an *open* action to the object referenced |  
 28778 by `file_actions` that shall cause the file named by `path` to be opened (as if `open(path, oflag, mode)` |  
 28779 had been called, and the returned file descriptor, if not `fildes`, had been changed to `fildes`) when a |  
 28780 new process is spawned using this file actions object. If `fildes` was already an open file descriptor, |  
 28781 it shall be closed before the new file is opened.

28782 The string described by `path` shall be copied by the `posix_spawn_file_actions_addopen()` function. |

28783 **RETURN VALUE**

28784 Upon successful completion, these functions shall return zero; otherwise, an error number shall  
 28785 be returned to indicate the error.

28786 **ERRORS**

28787 These functions shall fail if:

28788 [EBADF] The value specified by `fildes` is negative or greater than or equal to  
 28789 {OPEN\_MAX}.

28790 These functions may fail if:

28791 [EINVAL] The value specified by `file_actions` is invalid.

28792 [ENOMEM] Insufficient memory exists to add to the spawn file actions object.

28793 It shall not be considered an error for the *fildev* argument passed to these functions to specify a  
 28794 file descriptor for which the specified operation could not be performed at the time of the call.  
 28795 Any such error will be detected when the associated file actions object is later used during a  
 28796 *posix\_spawn()* or *posix\_spawnnp()* operation.

28797 **EXAMPLES**

28798 None.

28799 **APPLICATION USAGE**

28800 These functions are part of the Spawn option and need not be provided on all implementations.

28801 **RATIONALE**

28802 A spawn file actions object may be initialized to contain an ordered sequence of *close()*, *dup2()*,  
 28803 and *open()* operations to be used by *posix\_spawn()* or *posix\_spawnnp()* to arrive at the set of open  
 28804 file descriptors inherited by the spawned process from the set of open file descriptors in the  
 28805 parent at the time of the *posix\_spawn()* or *posix\_spawnnp()* call. It had been suggested that the  
 28806 *close()* and *dup2()* operations alone are sufficient to rearrange file descriptors, and that files  
 28807 which need to be opened for use by the spawned process can be handled either by having the  
 28808 calling process open them before the *posix\_spawn()* or *posix\_spawnnp()* call (and close them after),  
 28809 or by passing filenames to the spawned process (in *argv*) so that it may open them itself. The  
 28810 standard developers recommend that applications use one of these two methods when practical,  
 28811 since detailed error status on a failed open operation is always available to the application this  
 28812 way. However, the standard developers feel that allowing a spawn file actions object to specify  
 28813 open operations is still appropriate because:

- 28814 1. It is consistent with equivalent POSIX.5 (Ada) functionality.
- 28815 2. It supports the I/O redirection paradigm commonly employed by POSIX programs  
 28816 designed to be invoked from a shell. When such a program is the child process, it may not  
 28817 be designed to open files on its own.
- 28818 3. It allows file opens that might otherwise fail or violate file ownership/access rights if  
 28819 executed by the parent process.

28820 Regarding 2. above, note that the spawn open file action provides to *posix\_spawn()* and  
 28821 *posix\_spawnnp()* the same capability that the shell redirection operators provide to *system()*, only  
 28822 without the intervening execution of a shell; for example:

```
28823 system ("myprog <file1 3<file2");
```

28824 Regarding 3. above, note that if the calling process needs to open one or more files for access by  
 28825 the spawned process, but has insufficient spare file descriptors, then the open action is necessary  
 28826 to allow the *open()* to occur in the context of the child process after other file descriptors have  
 28827 been closed (that must remain open in the parent).

28828 Additionally, if a parent is executed from a file having a “set-user-id” mode bit set and the  
 28829 POSIX\_SPAWN\_RESETEUIDS flag is set in the spawn attributes, a file created within the parent  
 28830 process will (possibly incorrectly) have the parent’s effective user ID as its owner, whereas a file  
 28831 created via an *open()* action during *posix\_spawn()* or *posix\_spawnnp()* will have the parent’s real  
 28832 ID as its owner; and an open by the parent process may successfully open a file to which the real  
 28833 user should not have access or fail to open a file to which the real user should have access.

28834 **File Descriptor Mapping**

28835 The standard developers had originally proposed using an array which specified the mapping of  
28836 child file descriptors back to those of the parent. It was pointed out by the ballot group that it is  
28837 not possible to reshuffle file descriptors arbitrarily in a library implementation of *posix\_spawn()*  
28838 or *posix\_spawnnp()* without provision for one or more spare file descriptor entries (which simply  
28839 may not be available). Such an array requires that an implementation develop a complex  
28840 strategy to achieve the desired mapping without inadvertently closing the wrong file descriptor  
28841 at the wrong time.

28842 It was noted by a member of the Ada Language Bindings working group that the approved Ada  
28843 Language *Start\_Process* family of POSIX process primitives use a caller-specified set of file  
28844 actions to alter the normal *fork()/exec* semantics for inheritance of file descriptors in a very  
28845 flexible way, yet no such problems exist because the burden of determining how to achieve the  
28846 final file descriptor mapping is completely on the application. Furthermore, although the file  
28847 actions interface appears frightening at first glance, it is actually quite simple to implement in  
28848 either a library or the kernel.

28849 **FUTURE DIRECTIONS**

28850 None.

28851 **SEE ALSO**

28852 *close()*, *dup()*, *open()*, *posix\_spawn()*, *posix\_spawn\_file\_actions\_adddup2()*,  
28853 *posix\_spawn\_file\_actions\_destroy()*, *posix\_spawnnp()*, the Base Definitions volume of  
28854 IEEE Std 1003.1-200x, <spawn.h>

28855 **CHANGE HISTORY**

28856 First released in Issue 6. Derived from IEEE Std 1003.1d-1999.

28857 IEEE PASC Interpretation 1003.1 #105 is applied, adding a note to the DESCRIPTION that the  
28858 string pointed to by *path* is copied by the *posix\_spawn\_file\_actions\_addopen()* function.

28859 **NAME**

28860 posix\_spawn\_file\_actions\_adddup2 — add dup2 action to spawn file actions object  
 28861 (ADVANCED REALTIME)

28862 **SYNOPSIS**

```
28863 SPN #include <spawn.h>
28864
28864 int posix_spawn_file_actions_adddup2(posix_spawn_file_actions_t *
28865     file_actions, int fildes, int newfildes);
28866
```

28867 **DESCRIPTION**

28868 The *posix\_spawn\_file\_actions\_adddup2()* function shall add a *dup2()* action to the object |  
 28869 referenced by *file\_actions* that shall cause the file descriptor *fildes* to be duplicated as *newfildes* (as |  
 28870 if *dup2(fildes, newfildes)* had been called) when a new process is spawned using this file actions |  
 28871 object. |

28872 A spawn file actions object is as defined in *posix\_spawn\_file\_actions\_addclose()*. |

28873 **RETURN VALUE**

28874 Upon successful completion, the *posix\_spawn\_file\_actions\_adddup2()* function shall return zero;  
 28875 otherwise, an error number shall be returned to indicate the error.

28876 **ERRORS**

28877 The *posix\_spawn\_file\_actions\_adddup2()* function shall fail if:

28878 [EBADF] The value specified by *fildes* or *newfildes* is negative or greater than or equal to  
 28879 {OPEN\_MAX}.

28880 [ENOMEM] Insufficient memory exists to add to the spawn file actions object.

28881 The *posix\_spawn\_file\_actions\_adddup2()* function may fail if:

28882 [EINVAL] The value specified by *file\_actions* is invalid.

28883 It shall not be considered an error for the *fildes* argument passed to the  
 28884 *posix\_spawn\_file\_actions\_adddup2()* function to specify a file descriptor for which the specified  
 28885 operation could not be performed at the time of the call. Any such error will be detected when  
 28886 the associated file actions object is later used during a *posix\_spawn()* or *posix\_spawnnp()*  
 28887 operation.

28888 **EXAMPLES**

28889 None.

28890 **APPLICATION USAGE**

28891 The *posix\_spawn\_file\_actions\_adddup2()* function is part of the Spawn option and need not be  
 28892 provided on all implementations.

28893 **RATIONALE**

28894 Refer to the RATIONALE in *posix\_spawn\_file\_actions\_addclose()*.

28895 **FUTURE DIRECTIONS**

28896 None.

28897 **SEE ALSO**

28898 *dup()*, *posix\_spawn()*, *posix\_spawn\_file\_actions\_addclose()*, *posix\_spawn\_file\_actions\_destroy()*,  
 28899 *posix\_spawnnp()*, the Base Definitions volume of IEEE Std 1003.1-200x, <spawn.h>

28900 **CHANGE HISTORY**

28901 First released in Issue 6. Derived from IEEE Std 1003.1d-1999.

28902 IEEE PASC Interpretation 1003.1 #104 is applied, noting that the [EBADF] error can apply to the |

28903 *newfildes* argument in addition to *fildes*.

28904 **NAME**

28905 posix\_spawn\_file\_actions\_addopen — add open action to spawn file actions object  
28906 (ADVANCED REALTIME)

28907 **SYNOPSIS**

```
28908 SPN #include <spawn.h>  
28909 int posix_spawn_file_actions_addopen(posix_spawn_file_actions_t *restrict  
28910     file_actions, int fildes, const char *restrict path,  
28911     int oflag, mode_t mode);  
28912
```

28913 **DESCRIPTION**

28914 Refer to *posix\_spawn\_file\_actions\_addclose()*.

28915 **NAME**

28916 posix\_spawn\_file\_actions\_destroy, posix\_spawn\_file\_actions\_init — destroy and initialize  
 28917 spawn file actions object (**ADVANCED REALTIME**)

28918 **SYNOPSIS**

```
28919 SPN    #include <spawn.h>
28920
28921    int posix_spawn_file_actions_destroy(posix_spawn_file_actions_t *
28922        file_actions);
28923    int posix_spawn_file_actions_init(posix_spawn_file_actions_t *
28924        file_actions);
```

28925 **DESCRIPTION**

28926 The *posix\_spawn\_file\_actions\_destroy()* function shall destroy the object referenced by *file\_actions*; |  
 28927 the object becomes, in effect, uninitialized. An implementation may cause  
 28928 *posix\_spawn\_file\_actions\_destroy()* to set the object referenced by *file\_actions* to an invalid value. A  
 28929 destroyed spawn file actions object can be reinitialized using *posix\_spawn\_file\_actions\_init()*; the  
 28930 results of otherwise referencing the object after it has been destroyed are undefined.

28931 The *posix\_spawn\_file\_actions\_init()* function shall initialize the object referenced by *file\_actions* to  
 28932 contain no file actions for *posix\_spawn()* or *posix\_spawnnp()* to perform.

28933 A spawn file actions object is as defined in *posix\_spawn\_file\_actions\_addclose()*. |

28934 The effect of initializing an already initialized spawn file actions object is undefined. |

28935 **RETURN VALUE**

28936 Upon successful completion, these functions shall return zero; otherwise, an error number shall  
 28937 be returned to indicate the error.

28938 **ERRORS**

28939 The *posix\_spawn\_file\_actions\_init()* function shall fail if:

28940 [ENOMEM] Insufficient memory exists to initialize the spawn file actions object.

28941 The *posix\_spawn\_file\_actions\_destroy()* function may fail if:

28942 [EINVAL] The value specified by *file\_actions* is invalid.

28943 **EXAMPLES**

28944 None.

28945 **APPLICATION USAGE**

28946 These functions are part of the Spawn option and need not be provided on all implementations.

28947 **RATIONALE**

28948 Refer to the RATIONALE in *posix\_spawn\_file\_actions\_addclose()*.

28949 **FUTURE DIRECTIONS**

28950 None.

28951 **SEE ALSO**

28952 *posix\_spawn()*, *posix\_spawnnp()*, the Base Definitions volume of IEEE Std 1003.1-200x, **<spawn.h>**

28953 **CHANGE HISTORY**

28954 First released in Issue 6. Derived from IEEE Std 1003.1d-1999.

28955 In the SYNOPSIS, the inclusion of **<sys/types.h>** is no longer required.

28956 **NAME**

28957        `posix_spawn_file_actions_init` — initialize spawn file actions object (**ADVANCED REALTIME**)

28958 **SYNOPSIS**

28959 SPN        `#include <spawn.h>`

28960        `int posix_spawn_file_actions_init(posix_spawn_file_actions_t *`  
28961            `file_actions);`

28962

28963 **DESCRIPTION**

28964        Refer to `posix_spawn_file_actions_destroy()`.



28965 **NAME**

28966 `posix_spawnattr_destroy`, `posix_spawnattr_init` — destroy and initialize spawn attributes object  
 28967 (**ADVANCED REALTIME**)

28968 **SYNOPSIS**

28969 SPN `#include <spawn.h>`

28970 `int posix_spawnattr_destroy(posix_spawnattr_t *attr);`

28971 `int posix_spawnattr_init(posix_spawnattr_t *attr);`

28972

28973 **DESCRIPTION**

28974 The `posix_spawnattr_destroy()` function shall destroy a spawn attributes object. A destroyed `attr` |  
 28975 attributes object can be reinitialized using `posix_spawnattr_init()`; the results of otherwise |  
 28976 referencing the object after it has been destroyed are undefined. An implementation may cause |  
 28977 `posix_spawnattr_destroy()` to set the object referenced by `attr` to an invalid value.

28978 The `posix_spawnattr_init()` function shall initialize a spawn attributes object `attr` with the default |  
 28979 value for all of the individual attributes used by the implementation. Results are undefined if |  
 28980 `posix_spawnattr_init()` is called specifying an already initialized `attr` attributes object.

28981 A spawn attributes object is of type **posix\_spawnattr\_t** (defined in `<spawn.h>`) and is used to |  
 28982 specify the inheritance of process attributes across a spawn operation. IEEE Std 1003.1-200x does |  
 28983 not define comparison or assignment operators for the type **posix\_spawnattr\_t**.

28984 Each implementation shall document the individual attributes it uses and their default values |  
 28985 unless these values are defined by IEEE Std 1003.1-200x. Attributes not defined by |  
 28986 IEEE Std 1003.1-200x, their default values, and the names of the associated functions to get and |  
 28987 set those attribute values are implementation-defined.

28988 The resulting spawn attributes object (possibly modified by setting individual attribute values), |  
 28989 is used to modify the behavior of `posix_spawn()` or `posix_spawnp()`. After a spawn attributes |  
 28990 object has been used to spawn a process by a call to a `posix_spawn()` or `posix_spawnp()`, any |  
 28991 function affecting the attributes object (including destruction) shall not affect any process that |  
 28992 has been spawned in this way.

28993 **RETURN VALUE**

28994 Upon successful completion, `posix_spawnattr_destroy()` and `posix_spawnattr_init()` shall return |  
 28995 zero; otherwise, an error number shall be returned to indicate the error.

28996 **ERRORS**

28997 The `posix_spawnattr_init()` function shall fail if:

28998 [ENOMEM] Insufficient memory exists to initialize the spawn attributes object.

28999 The `posix_spawnattr_destroy()` function may fail if:

29000 [EINVAL] The value specified by `attr` is invalid.

29001 **EXAMPLES**

29002 None.

29003 **APPLICATION USAGE**

29004 These functions are part of the Spawn option and need not be provided on all implementations.

29005 **RATIONALE**

29006 The original spawn interface proposed in IEEE Std 1003.1-200x defined the attributes that specify |  
 29007 the inheritance of process attributes across a spawn operation as a structure. In order to be able |  
 29008 to separate optional individual attributes under their appropriate options (that is, the `spawn-` |  
 29009 `schedparam` and `spawn-schedpolicy` attributes depending upon the Process Scheduling option), and

29010 also for extensibility and consistency with the newer POSIX interfaces, the attributes interface  
29011 has been changed to an opaque data type. This interface now consists of the type  
29012 **posix\_spawnattr\_t**, representing a spawn attributes object, together with associated functions to  
29013 initialize or destroy the attributes object, and to set or get each individual attribute. Although the  
29014 new object-oriented interface is more verbose than the original structure, it is simple to use,  
29015 more extensible, and easy to implement.

## 29016 FUTURE DIRECTIONS

29017 None.

## 29018 SEE ALSO

29019 *posix\_spawn()*, *posix\_spawnattr\_getsigdefault()*, *posix\_spawnattr\_getflags()*,  
29020 *posix\_spawnattr\_getpgroup()*, *posix\_spawnattr\_getschedparam()*, *posix\_spawnattr\_getschedpolicy()*,  
29021 *posix\_spawnattr\_getsigmask()*, *posix\_spawnattr\_setsigdefault()*, *posix\_spawnattr\_setflags()*,  
29022 *posix\_spawnattr\_setpgroup()*, *posix\_spawnattr\_setsigmask()*, *posix\_spawnattr\_setschedpolicy()*,  
29023 *posix\_spawnattr\_setschedparam()*, *posix\_spawn()*, the Base Definitions volume of  
29024 IEEE Std 1003.1-200x, <spawn.h>

## 29025 CHANGE HISTORY

29026 First released in Issue 6. Derived from IEEE Std 1003.1d-1999.

29027 IEEE PASC Interpretation 1003.1 #106 is applied, noting that the effect of initializing an already |  
29028 initialized spawn attributes option is undefined.

29029 **NAME**

29030 `posix_spawnattr_getflags`, `posix_spawnattr_setflags` — get and set spawn-flags attribute of  
 29031 spawn attributes object (**ADVANCED REALTIME**)

29032 **SYNOPSIS**

```
29033 SPN  #include <spawn.h>
29034      int posix_spawnattr_getflags(const posix_spawnattr_t *restrict attr,
29035                                 short *restrict flags);
29036      int posix_spawnattr_setflags(posix_spawnattr_t *attr, short flags);
29037
```

29038 **DESCRIPTION**

29039 The `posix_spawnattr_getflags()` function shall obtain the value of the *spawn-flags* attribute from  
 29040 the attributes object referenced by *attr*.

29041 The `posix_spawnattr_setflags()` function shall set the *spawn-flags* attribute in an initialized  
 29042 attributes object referenced by *attr*.

29043 The *spawn-flags* attribute is used to indicate which process attributes are to be changed in the  
 29044 new process image when invoking `posix_spawn()` or `posix_spawnp()`. It is the bitwise-inclusive  
 29045 OR of zero or more of the following flags:

```
29046 POSIX_SPAWN_RESETEIDS
29047 POSIX_SPAWN_SETPGROUP
29048 POSIX_SPAWN_SETSIGDEF
29049 POSIX_SPAWN_SETSIGMASK
29050 PS  POSIX_SPAWN_SETSCHEDPARAM
29051    POSIX_SPAWN_SETSCHEDULER
29052
```

29053 These flags are defined in `<spawn.h>`. The default value of this attribute shall be as if no flags  
 29054 were set.

29055 **RETURN VALUE**

29056 Upon successful completion, `posix_spawnattr_getflags()` shall return zero and store the value of  
 29057 the *spawn-flags* attribute of *attr* into the object referenced by the *flags* parameter; otherwise, an  
 29058 error number shall be returned to indicate the error.

29059 Upon successful completion, `posix_spawnattr_setflags()` shall return zero; otherwise, an error  
 29060 number shall be returned to indicate the error.

29061 **ERRORS**

29062 These functions may fail if:

29063 [EINVAL] The value specified by *attr* is invalid.

29064 The `posix_spawnattr_setflags()` function may fail if:

29065 [EINVAL] The value of the attribute being set is not valid.

29066 **EXAMPLES**

29067           None.

29068 **APPLICATION USAGE**

29069           These functions are part of the Spawn option and need not be provided on all implementations.

29070 **RATIONALE**

29071           None.

29072 **FUTURE DIRECTIONS**

29073           None.

29074 **SEE ALSO**

29075           *posix\_spawn()*, *posix\_spawnattr\_destroy()*, *posix\_spawnattr\_init()*, *posix\_spawnattr\_getsigdefault()*,  
29076           *posix\_spawnattr\_getpgroup()*, *posix\_spawnattr\_getschedparam()*, *posix\_spawnattr\_getschedpolicy()*,  
29077           *posix\_spawnattr\_getsigmask()*, *posix\_spawnattr\_setsigdefault()*, *posix\_spawnattr\_setpgroup()*,  
29078           *posix\_spawnattr\_setschedparam()*, *posix\_spawnattr\_setschedpolicy()*, *posix\_spawnattr\_setsigmask()*,  
29079           *posix\_spawnnp()*, the Base Definitions volume of IEEE Std 1003.1-200x, <**spawn.h**>

29080 **CHANGE HISTORY**

29081           First released in Issue 6. Derived from IEEE Std 1003.1d-1999.

29082 **NAME**

29083 posix\_spawnattr\_getpgroup, posix\_spawnattr\_setpgroup — get and set spawn-pgroup attribute  
 29084 of spawn attributes object (**ADVANCED REALTIME**)

29085 **SYNOPSIS**

29086 SPN `#include <spawn.h>`

29087 `int posix_spawnattr_getpgroup(const posix_spawnattr_t *restrict attr,`  
 29088 `pid_t *restrict pgroup);`

29089 `int posix_spawnattr_setpgroup(posix_spawnattr_t *attr, pid_t pgroup);`  
 29090

29091 **DESCRIPTION**

29092 The *posix\_spawnattr\_getpgroup()* function shall obtain the value of the *spawn-pgroup* attribute |  
 29093 from the attributes object referenced by *attr*.

29094 The *posix\_spawnattr\_setpgroup()* function shall set the *spawn-pgroup* attribute in an initialized |  
 29095 attributes object referenced by *attr*.

29096 The *spawn-pgroup* attribute represents the process group to be joined by the new process image |  
 29097 in a spawn operation (if `POSIX_SPAWN_SETPGROUP` is set in the *spawn-flags* attribute). The |  
 29098 default value of this attribute shall be zero. |

29099 **RETURN VALUE**

29100 Upon successful completion, *posix\_spawnattr\_getpgroup()* shall return zero and store the value of  
 29101 the *spawn-pgroup* attribute of *attr* into the object referenced by the *pgroup* parameter; otherwise,  
 29102 an error number shall be returned to indicate the error.

29103 Upon successful completion, *posix\_spawnattr\_setpgroup()* shall return zero; otherwise, an error  
 29104 number shall be returned to indicate the error.

29105 **ERRORS**

29106 These functions may fail if:

29107 [EINVAL] The value specified by *attr* is invalid.

29108 The *posix\_spawnattr\_setpgroup()* function may fail if:

29109 [EINVAL] The value of the attribute being set is not valid.

29110 **EXAMPLES**

29111 None.

29112 **APPLICATION USAGE**

29113 These functions are part of the Spawn option and need not be provided on all implementations.

29114 **RATIONALE**

29115 None.

29116 **FUTURE DIRECTIONS**

29117 None.

29118 **SEE ALSO**

29119 *posix\_spawn()*, *posix\_spawnattr\_destroy()*, *posix\_spawnattr\_init()*, *posix\_spawnattr\_getsigdefault()*,  
 29120 *posix\_spawnattr\_getflags()*, *posix\_spawnattr\_getschedparam()*, *posix\_spawnattr\_getschedpolicy()*,  
 29121 *posix\_spawnattr\_getsigmask()*, *posix\_spawnattr\_setsigdefault()*, *posix\_spawnattr\_setflags()*,  
 29122 *posix\_spawnattr\_setschedparam()*, *posix\_spawnattr\_setschedpolicy()*, *posix\_spawnattr\_setsigmask()*,  
 29123 *posix\_spawnnp()*, the Base Definitions volume of IEEE Std 1003.1-200x, `<spawn.h>`

29124 **CHANGE HISTORY**

29125 First released in Issue 6. Derived from IEEE Std 1003.1d-1999.

29126 **NAME**

29127 `posix_spawnattr_getschedparam`, `posix_spawnattr_setschedparam` — get and set spawn-  
 29128 `schedparam` attribute of spawn attributes object (**ADVANCED REALTIME**)

29129 **SYNOPSIS**

29130 SPN PS `#include <spawn.h>`

29131 `#include <sched.h>`

```
29132 int posix_spawnattr_getschedparam(
29133     const posix_spawnattr_t *restrict attr,
29134     struct sched_param *restrict schedparam);
29135 int posix_spawnattr_setschedparam(posix_spawnattr_t *restrict attr,
29136     const struct sched_param *restrict schedparam);
29137
```

29138 **DESCRIPTION**

29139 The `posix_spawnattr_getschedparam()` function shall obtain the value of the `spawn-schedparam` |  
 29140 attribute from the attributes object referenced by `attr`.

29141 The `posix_spawnattr_setschedparam()` function shall set the `spawn-schedparam` attribute in an |  
 29142 initialized attributes object referenced by `attr`.

29143 The `spawn-schedparam` attribute represents the scheduling parameters to be assigned to the new |  
 29144 process image in a spawn operation (if `POSIX_SPAWN_SETSCHEDULER` or |  
 29145 `POSIX_SPAWN_SETSCHEDPARAM` is set in the `spawn-flags` attribute). The default value of this |  
 29146 attribute is unspecified.

29147 **RETURN VALUE**

29148 Upon successful completion, `posix_spawnattr_getschedparam()` shall return zero and store the  
 29149 value of the `spawn-schedparam` attribute of `attr` into the object referenced by the `schedparam`  
 29150 parameter; otherwise, an error number shall be returned to indicate the error.

29151 Upon successful completion, `posix_spawnattr_setschedparam()` shall return zero; otherwise, an  
 29152 error number shall be returned to indicate the error.

29153 **ERRORS**

29154 These functions may fail if:

29155 [EINVAL] The value specified by `attr` is invalid.

29156 The `posix_spawnattr_setschedparam()` function may fail if:

29157 [EINVAL] The value of the attribute being set is not valid.

29158 **EXAMPLES**

29159 None.

29160 **APPLICATION USAGE**

29161 These functions are part of the Spawn and Process Scheduling options and need not be provided  
 29162 on all implementations.

29163 **RATIONALE**

29164 None.

29165 **FUTURE DIRECTIONS**

29166 None.

29167 **SEE ALSO**

29168 *posix\_spawn()*, *posix\_spawnattr\_destroy()*, *posix\_spawnattr\_init()*, *posix\_spawnattr\_getsigdefault()*,  
29169 *posix\_spawnattr\_getflags()*, *posix\_spawnattr\_getpgroup()*, *posix\_spawnattr\_getschedpolicy()*,  
29170 *posix\_spawnattr\_getsigmask()*, *posix\_spawnattr\_setsigdefault()*, *posix\_spawnattr\_setflags()*,  
29171 *posix\_spawnattr\_setpgroup()*, *posix\_spawnattr\_setschedpolicy()*, *posix\_spawnattr\_setsigmask()*,  
29172 *posix\_spawnp()*, the Base Definitions volume of IEEE Std 1003.1-200x, <sched.h>, <spawn.h>

29173 **CHANGE HISTORY**

29174 First released in Issue 6. Derived from IEEE Std 1003.1d-1999.



29175 **NAME**

29176 `posix_spawnattr_getschedpolicy`, `posix_spawnattr_setschedpolicy` — get and set spawn-  
 29177 `schedpolicy` attribute of spawn attributes object (**ADVANCED REALTIME**)

29178 **SYNOPSIS**

```
29179 SPN PS #include <spawn.h>
```

```
29180 #include <sched.h>
```

```
29181 int posix_spawnattr_getschedpolicy(
```

```
29182     const posix_spawnattr_t *restrict attr,
```

```
29183     int *restrict schedpolicy);
```

```
29184 int posix_spawnattr_setschedpolicy(posix_spawnattr_t *attr,
```

```
29185     int schedpolicy);
```

```
29186
```

29187 **DESCRIPTION**

29188 The `posix_spawnattr_getschedpolicy()` function shall obtain the value of the *spawn-schedpolicy* |  
 29189 attribute from the attributes object referenced by *attr*.

29190 The `posix_spawnattr_setschedpolicy()` function shall set the *spawn-schedpolicy* attribute in an |  
 29191 initialized attributes object referenced by *attr*.

29192 The *spawn-schedpolicy* attribute represents the scheduling policy to be assigned to the new |  
 29193 process image in a spawn operation (if `POSIX_SPAWN_SETSCHEDULER` is set in the *spawn-* |  
 29194 *flags* attribute). The default value of this attribute is unspecified.

29195 **RETURN VALUE**

29196 Upon successful completion, `posix_spawnattr_getschedpolicy()` shall return zero and store the  
 29197 value of the *spawn-schedpolicy* attribute of *attr* into the object referenced by the *schedpolicy*  
 29198 parameter; otherwise, an error number shall be returned to indicate the error.

29199 Upon successful completion, `posix_spawnattr_setschedpolicy()` shall return zero; otherwise, an  
 29200 error number shall be returned to indicate the error.

29201 **ERRORS**

29202 These functions may fail if:

29203 [EINVAL] The value specified by *attr* is invalid.

29204 The `posix_spawnattr_setschedpolicy()` function may fail if:

29205 [EINVAL] The value of the attribute being set is not valid.

29206 **EXAMPLES**

29207 None.

29208 **APPLICATION USAGE**

29209 These functions are part of the Spawn and Process Scheduling options and need not be provided  
 29210 on all implementations.

29211 **RATIONALE**

29212 None.

29213 **FUTURE DIRECTIONS**

29214 None.

29215 **SEE ALSO**

29216 `posix_spawn()`, `posix_spawnattr_destroy()`, `posix_spawnattr_init()`, `posix_spawnattr_getsigdefault()`,

29217 `posix_spawnattr_getflags()`, `posix_spawnattr_getpgroup()`, `posix_spawnattr_getschedparam()`,

29218 `posix_spawnattr_getsigmask()`, `posix_spawnattr_setsigdefault()`, `posix_spawnattr_setflags()`,

29219            *posix\_spawnattr\_setpgroup()*, *posix\_spawnattr\_setschedparam()*, *posix\_spawnattr\_setsigmask()*,  
29220            *posix\_spawnnp()*, the Base Definitions volume of IEEE Std 1003.1-200x, <sched.h>, <spawn.h>

29221 **CHANGE HISTORY**

29222            First released in Issue 6. Derived from IEEE Std 1003.1d-1999.

29223 **NAME**

29224 `posix_spawnattr_getsigdefault`, `posix_spawnattr_setsigdefault` — get and set spawn-sigdefault  
 29225 attribute of spawn attributes object (**ADVANCED REALTIME**)

29226 **SYNOPSIS**

```
29227 SPN #include <signal.h>
29228 #include <spawn.h>

29229 int posix_spawnattr_getsigdefault(
29230     const posix_spawnattr_t *restrict attr,
29231     sigset_t *restrict sigdefault);
29232 int posix_spawnattr_setsigdefault(posix_spawnattr_t *restrict attr,
29233     const sigset_t *restrict sigdefault);
29234
```

29235 **DESCRIPTION**

29236 The `posix_spawnattr_getsigdefault()` function shall obtain the value of the `spawn-sigdefault` |  
 29237 attribute from the attributes object referenced by `attr`.

29238 The `posix_spawnattr_setsigdefault()` function shall set the `spawn-sigdefault` attribute in an |  
 29239 initialized attributes object referenced by `attr`.

29240 The `spawn-sigdefault` attribute represents the set of signals to be forced to default signal handling |  
 29241 in the new process image (if `POSIX_SPAWN_SETSIGDEF` is set in the `spawn-flags` attribute) by a |  
 29242 spawn operation. The default value of this attribute shall be an empty signal set. |

29243 **RETURN VALUE**

29244 Upon successful completion, `posix_spawnattr_getsigdefault()` shall return zero and store the value  
 29245 of the `spawn-sigdefault` attribute of `attr` into the object referenced by the `sigdefault` parameter;  
 29246 otherwise, an error number shall be returned to indicate the error.

29247 Upon successful completion, `posix_spawnattr_setsigdefault()` shall return zero; otherwise, an error  
 29248 number shall be returned to indicate the error.

29249 **ERRORS**

29250 These functions may fail if:

29251 [EINVAL] The value specified by `attr` is invalid.

29252 The `posix_spawnattr_setsigdefault()` function may fail if:

29253 [EINVAL] The value of the attribute being set is not valid.

29254 **EXAMPLES**

29255 None.

29256 **APPLICATION USAGE**

29257 These functions are part of the Spawn option and need not be provided on all implementations.

29258 **RATIONALE**

29259 None.

29260 **FUTURE DIRECTIONS**

29261 None.

29262 **SEE ALSO**

29263 `posix_spawn()`, `posix_spawnattr_destroy()`, `posix_spawnattr_init()`, `posix_spawnattr_getflags()`,  
 29264 `posix_spawnattr_getpgroup()`, `posix_spawnattr_getschedparam()`, `posix_spawnattr_getschedpolicy()`,  
 29265 `posix_spawnattr_getsigmask()`, `posix_spawnattr_setflags()`, `posix_spawnattr_setpgroup()`,  
 29266 `posix_spawnattr_setschedparam()`, `posix_spawnattr_setschedpolicy()`, `posix_spawnattr_setsigmask()`,

- 29267            *posix\_spawnp()*, the Base Definitions volume of IEEE Std 1003.1-200x, <signal.h>, <spawn.h>
- 29268 **CHANGE HISTORY**
- 29269            First released in Issue 6. Derived from IEEE Std 1003.1d-1999.

29270 **NAME**

29271 posix\_spawnattr\_getsigmask, posix\_spawnattr\_setsigmask — get and set spawn-sigmask  
 29272 attribute of spawn attributes object (**ADVANCED REALTIME**)

29273 **SYNOPSIS**

```
29274 SPN #include <signal.h>
29275 #include <spawn.h>

29276 int posix_spawnattr_getsigmask(const posix_spawnattr_t *restrict attr,
29277 sigset_t *restrict sigmask);
29278 int posix_spawnattr_setsigmask(posix_spawnattr_t *restrict attr,
29279 const sigset_t *restrict sigmask);
29280
```

29281 **DESCRIPTION**

29282 The *posix\_spawnattr\_getsigmask()* function shall obtain the value of the *spawn-sigmask* attribute  
 29283 from the attributes object referenced by *attr*.

29284 The *posix\_spawnattr\_setsigmask()* function shall set the *spawn-sigmask* attribute in an initialized  
 29285 attributes object referenced by *attr*.

29286 The *spawn-sigmask* attribute represents the signal mask in effect in the new process image of a  
 29287 spawn operation (if `POSIX_SPAWN_SETSIGMASK` is set in the *spawn-flags* attribute). The  
 29288 default value of this attribute is unspecified.

29289 **RETURN VALUE**

29290 Upon successful completion, *posix\_spawnattr\_getsigmask()* shall return zero and store the value  
 29291 of the *spawn-sigmask* attribute of *attr* into the object referenced by the *sigmask* parameter;  
 29292 otherwise, an error number shall be returned to indicate the error.

29293 Upon successful completion, *posix\_spawnattr\_setsigmask()* shall return zero; otherwise, an error  
 29294 number shall be returned to indicate the error.

29295 **ERRORS**

29296 These functions may fail if:

29297 [EINVAL] The value specified by *attr* is invalid.

29298 The *posix\_spawnattr\_setsigmask()* function may fail if:

29299 [EINVAL] The value of the attribute being set is not valid.

29300 **EXAMPLES**

29301 None.

29302 **APPLICATION USAGE**

29303 These functions are part of the Spawn option and need not be provided on all implementations.

29304 **RATIONALE**

29305 None.

29306 **FUTURE DIRECTIONS**

29307 None.

29308 **SEE ALSO**

29309 *posix\_spawn()*, *posix\_spawnattr\_destroy()*, *posix\_spawnattr\_init()*, *posix\_spawnattr\_getsigdefault()*,  
 29310 *posix\_spawnattr\_getflags()*, *posix\_spawnattr\_getpgroup()*, *posix\_spawnattr\_getschedparam()*,  
 29311 *posix\_spawnattr\_getschedpolicy()*, *posix\_spawnattr\_setsigdefault()*, *posix\_spawnattr\_setflags()*,  
 29312 *posix\_spawnattr\_setpgroup()*, *posix\_spawnattr\_setschedparam()*, *posix\_spawnattr\_setschedpolicy()*,  
 29313 *posix\_spawnnp()*, the Base Definitions volume of IEEE Std 1003.1-200x, `<signal.h>`, `<spawn.h>`

29314 **CHANGE HISTORY**

29315 First released in Issue 6. Derived from IEEE Std 1003.1d-1999.

29316 **NAME**

29317 `posix_spawnattr_init` — initialize spawn attributes object (**ADVANCED REALTIME**)

29318 **SYNOPSIS**

29319 SPN `#include <spawn.h>`

29320 `int posix_spawnattr_init(posix_spawnattr_t *attr);`

29321

29322 **DESCRIPTION**

29323 Refer to `posix_spawnattr_destroy()`.

29324 **NAME**

29325        `posix_spawnattr_setflags` — set spawn-flags attribute of spawn attributes object (**ADVANCED**  
29326        **REALTIME**)

29327 **SYNOPSIS**

29328 SPN    `#include <spawn.h>`

29329        `int posix_spawnattr_setflags(posix_spawnattr_t *attr, short flags);`

29330

29331 **DESCRIPTION**

29332        Refer to `posix_spawnattr_getflags()`.



29333 **NAME**

29334 `posix_spawnattr_setpgroup` — set spawn-pgroup attribute of spawn attributes object  
29335 (**ADVANCED REALTIME**)

29336 **SYNOPSIS**

29337 SPN `#include <spawn.h>`

29338 `int posix_spawnattr_setpgroup(posix_spawnattr_t *attr, pid_t pgroup);`  
29339

29340 **DESCRIPTION**

29341 Refer to `posix_spawnattr_getpgroup()`.

29342 **NAME**

29343 posix\_spawnattr\_setschedparam — set spawn-schedparam attribute of spawn attributes object  
29344 (**ADVANCED REALTIME**)

29345 **SYNOPSIS**

29346 SPN PS #include <sched.h>

29347 #include <spawn.h>

```
29348 int posix_spawnattr_setschedparam(posix_spawnattr_t *restrict attr,  
29349     const struct sched_param *restrict schedparam);
```

29350

29351 **DESCRIPTION**

29352 Refer to *posix\_spawnattr\_getschedparam()*.

29353 **NAME**

29354 `posix_spawnattr_setschedpolicy` — set spawn-schedpolicy attribute of spawn attributes object  
29355 (**ADVANCED REALTIME**)

29356 **SYNOPSIS**

29357 SPN PS `#include <sched.h>`

29358 `#include <spawn.h>`

```
29359 int posix_spawnattr_setschedpolicy(posix_spawnattr_t *attr,  
29360 int schedpolicy);
```

29361

29362 **DESCRIPTION**

29363 Refer to `posix_spawnattr_getschedpolicy()`.

29364 **NAME**

29365        posix\_spawnattr\_setsigdefault — set spawn-sigdefault attribute of spawn attributes object  
29366        (ADVANCED REALTIME)

29367 **SYNOPSIS**

29368 SPN    #include <signal.h>

29369        #include <spawn.h>

```
29370        int posix_spawnattr_setsigdefault(posix_spawnattr_t *restrict attr,  
29371                                        const sigset_t *restrict sigdefault);
```

29372

29373 **DESCRIPTION**

29374        Refer to *posix\_spawnattr\_getsigdefault()*.

29375 **NAME**

29376 `posix_spawnattr_setsigmask` — set spawn-sigmask attribute of spawn attributes object  
29377 (**ADVANCED REALTIME**)

29378 **SYNOPSIS**

29379 SPN `#include <signal.h>`

29380 `#include <spawn.h>`

```
29381 int posix_spawnattr_setsigmask(posix_spawnattr_t *restrict attr,  
29382     const sigset_t *restrict sigmask);
```

29383

29384 **DESCRIPTION**

29385 Refer to `posix_spawnattr_getsigmask()`.

29386 **NAME**

29387 posix\_spawnnp — spawn a process (**ADVANCED REALTIME**)

29388 **SYNOPSIS**

29389 SPN `#include <spawn.h>`

```
29390 int posix_spawnnp(pid_t *restrict pid, const char *restrict file,  
29391                  const posix_spawn_file_actions_t *file_actions,  
29392                  const posix_spawnattr_t *restrict attrp,  
29393                  char *const argv[restrict], char *const envp[restrict]);  
29394
```

29395 **DESCRIPTION**

29396 Refer to *posix\_spawn()*.

29397 **NAME**

29398        posix\_trace\_attr\_destroy, posix\_trace\_attr\_init — trace stream attributes object destroy and  
 29399        initialization (**TRACING**)

29400 **SYNOPSIS**

```
29401 TRC    #include <trace.h>
```

```
29402        int posix_trace_attr_destroy(trace_attr_t *attr);
```

```
29403        int posix_trace_attr_init(trace_attr_t *attr);
```

29404

29405 **DESCRIPTION**

29406        The *posix\_trace\_attr\_destroy()* function shall destroy an initialized trace attributes object. A  
 29407        destroyed *attr* attributes object can be reinitialized using *posix\_trace\_attr\_init()*; the results of  
 29408        otherwise referencing the object after it has been destroyed are undefined.

29409        The *posix\_trace\_attr\_init()* function shall initialize a trace attributes object *attr* with the default  
 29410        value for all of the individual attributes used by a given implementation. The read-only  
 29411        *generation-version* and *clock-resolution* attributes of the newly initialized trace attributes object  
 29412        shall be set to their appropriate values (see Section 2.11.1.2 (on page 525)).

29413        Results are undefined if *posix\_trace\_attr\_init()* is called specifying an already initialized *attr*  
 29414        attributes object.

29415        Implementations may add extensions to the trace attributes object structure as permitted in the  
 29416        Base Definitions volume of IEEE Std 1003.1-200x, Chapter 2, Conformance.

29417        The resulting attributes object (possibly modified by setting individual attributes values), when  
 29418        used by *posix\_trace\_create()*, defines the attributes of the trace stream created. A single attributes  
 29419        object can be used in multiple calls to *posix\_trace\_create()*. After one or more trace streams have  
 29420        been created using an attributes object, any function affecting that attributes object, including  
 29421        destruction, shall not affect any trace stream previously created. An initialized attributes object  
 29422        also serves to receive the attributes of an existing trace stream or trace log when calling the  
 29423        *posix\_trace\_get\_attr()* function.

29424 **RETURN VALUE**

29425        Upon successful completion, these functions shall return a value of zero. Otherwise, they shall  
 29426        return the corresponding error number.

29427 **ERRORS**

29428        The *posix\_trace\_attr\_destroy()* function may fail if:

29429        [EINVAL]        The value of *attr* is invalid.

29430        The *posix\_trace\_attr\_init()* function shall fail if:

29431        [ENOMEM]       Insufficient memory exists to initialize the trace attributes object.

29432 **EXAMPLES**

29433        None.

29434 **APPLICATION USAGE**

29435        None.

29436 **RATIONALE**

29437        None.

29438 **FUTURE DIRECTIONS**

29439       None.

29440 **SEE ALSO**

29441       *posix\_trace\_create()*, *posix\_trace\_get\_attr()*, *uname()*, the Base Definitions volume of  
29442       IEEE Std 1003.1-200x, <**trace.h**>

29443 **CHANGE HISTORY**

29444       First released in Issue 6. Derived from IEEE Std 1003.1q-2000. |

29445       IEEE PASC Interpretation 1003.1 #123 is applied. |



29446 **NAME**

29447 `posix_trace_attr_getclockres`, `posix_trace_attr_getcreatetime`, `posix_trace_attr_getgenversion`,  
 29448 `posix_trace_attr_getname`, `posix_trace_attr_setname` — retrieve and set information about a  
 29449 trace stream (**TRACING**)

29450 **SYNOPSIS**

```
29451 TRC #include <time.h>
29452 #include <trace.h>

29453 int posix_trace_attr_getclockres(const trace_attr_t *attr,
29454     struct timespec *resolution);
29455 int posix_trace_attr_getcreatetime(const trace_attr_t *attr,
29456     struct timespec *createtime);

29457 #include <trace.h>

29458 int posix_trace_attr_getgenversion(const trace_attr_t *attr,
29459     char *genversion);
29460 int posix_trace_attr_getname(const trace_attr_t *attr,
29461     char *tracename);
29462 int posix_trace_attr_setname(trace_attr_t *attr,
29463     const char *tracename);
29464
```

29465 **DESCRIPTION**

29466 The `posix_trace_attr_getclockres()` function shall copy the clock resolution of the clock used to  
 29467 generate timestamps from the *clock-resolution* attribute of the attributes object pointed to by the  
 29468 *attr* argument into the structure pointed to by the *resolution* argument.

29469 The `posix_trace_attr_getcreatetime()` function shall copy the trace stream creation time from the  
 29470 *creation-time* attribute of the attributes object pointed to by the *attr* argument into the structure  
 29471 pointed to by the *createtime* argument. The *creation-time* attribute shall represent the time of  
 29472 creation of the trace stream.

29473 The `posix_trace_attr_getgenversion()` function shall copy the string containing version information  
 29474 from the *generation-version* attribute of the attributes object pointed to by the *attr* argument into  
 29475 the string pointed to by the *genversion* argument. The *genversion* argument shall be the address of  
 29476 a character array which can store at least {TRACE\_NAME\_MAX} characters.

29477 The `posix_trace_attr_getname()` function shall copy the string containing the trace name from the  
 29478 *trace-name* attribute of the attributes object pointed to by the *attr* argument into the string  
 29479 pointed to by the *tracename* argument. The *tracename* argument shall be the address of a character  
 29480 array which can store at least {TRACE\_NAME\_MAX} characters.

29481 The `posix_trace_attr_setname()` function shall set the name in the *trace-name* attribute of the  
 29482 attributes object pointed to by the *attr* argument, using the trace name string supplied by the  
 29483 *tracename* argument. If the supplied string contains more than {TRACE\_NAME\_MAX}  
 29484 characters, the name copied into the *trace-name* attribute may be truncated to one less than the  
 29485 length of {TRACE\_NAME\_MAX} characters. The default value is a null string.

29486 **RETURN VALUE**

29487 Upon successful completion, these functions shall return a value of zero. Otherwise, they shall  
 29488 return the corresponding error number.

29489 If successful, the `posix_trace_attr_getclockres()` function stores the *clock-resolution* attribute value  
 29490 in the object pointed to by *resolution*. Otherwise, the content of this object is unspecified.

29491 If successful, the *posix\_trace\_attr\_getcreatetime()* function stores the trace stream creation time in  
29492 the object pointed to by *createtime*. Otherwise, the content of this object is unspecified.

29493 If successful, the *posix\_trace\_attr\_getgenversion()* function stores the trace version information in  
29494 the string pointed to by *genversion*. Otherwise, the content of this string is unspecified.

29495 If successful, the *posix\_trace\_attr\_getname()* function stores the trace name in the string pointed  
29496 to by *tracename*. Otherwise, the content of this string is unspecified.

## 29497 ERRORS

29498 The *posix\_trace\_attr\_getclockres()*, *posix\_trace\_attr\_getcreatetime()*, *posix\_trace\_attr\_getgenversion()*,  
29499 and *posix\_trace\_attr\_getname()* functions may fail if:

29500 [EINVAL] The value specified by one of the arguments is invalid.

## 29501 EXAMPLES

29502 None.

## 29503 APPLICATION USAGE

29504 None.

## 29505 RATIONALE

29506 None.

## 29507 FUTURE DIRECTIONS

29508 None.

## 29509 SEE ALSO

29510 *posix\_trace\_attr\_init()*, *posix\_trace\_create()*, *posix\_trace\_get\_attr()*, *uname()*, the Base Definitions  
29511 volume of IEEE Std 1003.1-200x, <**time.h**>, <**trace.h**>

## 29512 CHANGE HISTORY

29513 First released in Issue 6. Derived from IEEE Std 1003.1q-2000.

29514 **NAME**

29515        posix\_trace\_attr\_getcreatetime — retrieve and set information about a trace stream (**TRACING**)

29516 **SYNOPSIS**

29517 TRC     #include <time.h>

29518        #include <trace.h>

```
29519        int posix_trace_attr_getcreatetime(const trace_attr_t *attr,  
29520                                          struct timespec *createtime);
```

29521

29522 **DESCRIPTION**

29523        Refer to *posix\_trace\_attr\_getclockres()*.

29524 **NAME**

29525        posix\_trace\_attr\_getgenversion — retrieve and set information about a trace stream  
29526        **(TRACING)**

29527 **SYNOPSIS**

29528 TRC    #include <trace.h>

```
29529        int posix_trace_attr_getgenversion(const trace_attr_t *attr,  
29530        char *genversion);
```

29531

29532 **DESCRIPTION**

29533        Refer to *posix\_trace\_attr\_getclockres()*.

## 29534 NAME

29535 posix\_trace\_attr\_getinherited, posix\_trace\_attr\_getlogfullpolicy,  
 29536 posix\_trace\_attr\_getstreamfullpolicy, posix\_trace\_attr\_setinherited,  
 29537 posix\_trace\_attr\_setlogfullpolicy, posix\_trace\_attr\_setstreamfullpolicy — retrieve and set the  
 29538 behavior of a trace stream (**TRACING**)

## 29539 SYNOPSIS

```
29540 TRC #include <trace.h>
29541 TRC TRI int posix_trace_attr_getinherited(const trace_attr_t *restrict attr,
29542 int *restrict inheritancepolicy);
29543 TRC TRL int posix_trace_attr_getlogfullpolicy(const trace_attr_t *restrict attr,
29544 int *restrict logpolicy);
29545 TRC int posix_trace_attr_getstreamfullpolicy(const trace_attr_t *attr,
29546 int *streampolicy);
29547 TRC TRI int posix_trace_attr_setinherited(trace_attr_t *attr,
29548 int inheritancepolicy);
29549 TRC TRL int posix_trace_attr_setlogfullpolicy(trace_attr_t *attr,
29550 int logpolicy);
29551 TRC int posix_trace_attr_setstreamfullpolicy(trace_attr_t *attr,
29552 int streampolicy);
29553
```

## 29554 DESCRIPTION

29555 TRI The *posix\_trace\_attr\_getinherited()* and *posix\_trace\_attr\_setinherited()* functions, respectively, shall  
 29556 get and set the inheritance policy stored in the *inheritance* attribute for traced processes across  
 29557 the *fork()* and *spawn()* operations. The *inheritance* attribute of the attributes object pointed to by  
 29558 the *attr* argument shall be set to one of the following values defined by manifest constants in the  
 29559 **<trace.h>** header:

## 29560 POSIX\_TRACE\_CLOSE\_FOR\_CHILD

29561 After a *fork()* or *spawn()* operation, the child shall not be traced, and tracing of the parent  
 29562 shall continue.

## 29563 POSIX\_TRACE\_INHERITED

29564 After a *fork()* or *spawn()* operation, if the parent is being traced, its child shall be  
 29565 concurrently traced using the same trace stream.

29566 The default value for the *inheritance* attribute is **POSIX\_TRACE\_CLOSE\_FOR\_CHILD**.

29567 TRL The *posix\_trace\_attr\_getlogfullpolicy()* and *posix\_trace\_attr\_setlogfullpolicy()* functions, |  
 29568 respectively, shall get and set the trace log full policy stored in the *log-full-policy* attribute of the |  
 29569 attributes object pointed to by the *attr* argument.

29570 The *log-full-policy* attribute shall be set to one of the following values defined by manifest  
 29571 constants in the **<trace.h>** header:

## 29572 POSIX\_TRACE\_LOOP

29573 The trace log shall loop until the associated trace stream is stopped. This policy means that  
 29574 when the trace log gets full, the file system shall reuse the resources allocated to the oldest  
 29575 trace events that were recorded. In this way, the trace log will always contain the most  
 29576 recent trace events flushed.

## 29577 POSIX\_TRACE\_UNTIL\_FULL

29578 The trace stream shall be flushed to the trace log until the trace log is full. This condition can  
 29579 be deduced from the *posix\_log\_full\_status* member status (see the *posix\_trace\_status\_info()*  
 29580 function). The last recorded trace event shall be the **POSIX\_TRACE\_STOP** trace event.

29581 POSIX\_TRACE\_APPEND  
 29582 The associated trace stream shall be flushed to the trace log without log size limitation. If  
 29583 the application specifies POSIX\_TRACE\_APPEND, the implementation shall ignore the  
 29584 *log-max-size* attribute.

29585 The default value for the *log-full-policy* attribute is POSIX\_TRACE\_LOOP.

29586 The *posix\_trace\_attr\_getstreamfullpolicy()* and *posix\_trace\_attr\_setstreamfullpolicy()* functions,  
 29587 respectively, shall get and set the trace stream full policy stored in the *stream-full-policy* attribute  
 29588 of the attributes object pointed to by the *attr* argument.

29589 The *stream-full-policy* attribute shall be set to one of the following values defined by manifest  
 29590 constants in the <trace.h> header:

29591 POSIX\_TRACE\_LOOP  
 29592 The trace stream shall loop until explicitly stopped by the *posix\_trace\_stop()* function. This  
 29593 policy means that when the trace stream is full, the trace system shall reuse the resources  
 29594 allocated to the oldest trace events recorded. In this way, the trace stream will always  
 29595 contain the most recent trace events recorded.

29596 POSIX\_TRACE\_UNTIL\_FULL  
 29597 The trace stream will run until the trace stream resources are exhausted. Then the trace  
 29598 stream will stop. This condition can be deduced from *posix\_stream\_status* and  
 29599 *posix\_stream\_full\_status* statuses (see the *posix\_trace\_status\_info()* function). When this trace  
 29600 stream is read, a POSIX\_TRACE\_STOP trace event shall be reported after reporting the last  
 29601 recorded trace event. The trace system shall reuse the resources allocated to any trace  
 29602 events already reported—see the *posix\_trace\_getnext\_event()*, *posix\_trace\_trygetnext\_event()*,  
 29603 and *posix\_trace\_timedgetnext\_event()* functions—or already flushed for an active trace stream  
 29604 with log if the Trace Log option is supported; see the *posix\_trace\_flush()* function. The trace  
 29605 system shall restart the trace stream when it is empty and may restart it sooner. A  
 29606 POSIX\_TRACE\_START trace event shall be reported before reporting the next recorded  
 29607 trace event.

29608 TRL POSIX\_TRACE\_FLUSH  
 29609 If the Trace Log option is supported, this policy is identical to the  
 29610 POSIX\_TRACE\_UNTIL\_FULL trace stream full policy except that the trace stream shall be  
 29611 flushed regularly as if *posix\_trace\_flush()* had been explicitly called. Defining this policy for  
 29612 an active trace stream without log shall be invalid.

29613 The default value for the *stream-full-policy* attribute shall be POSIX\_TRACE\_LOOP for an active  
 29614 trace stream without log.

29615 TRL If the Trace Log option is supported, the default value for the *stream-full-policy* attribute shall be  
 29616 POSIX\_TRACE\_FLUSH for an active trace stream with log.

29617 RETURN VALUE

29618 Upon successful completion, these functions shall return a value of zero. Otherwise, they shall  
 29619 return the corresponding error number.

29620 TRI If successful, the *posix\_trace\_attr\_getinherited()* function shall store the *inheritance* attribute value  
 29621 in the object pointed to by *inheritancepolicy*. Otherwise, the content of this object is undefined.

29622 TRL If successful, the *posix\_trace\_attr\_getlogfullpolicy()* function shall store the *log-full-policy* attribute  
 29623 value in the object pointed to by *logpolicy*. Otherwise, the content of this object is undefined.

29624 If successful, the *posix\_trace\_attr\_getstreamfullpolicy()* function shall store the *stream-full-policy*  
 29625 attribute value in the object pointed to by *streampolicy*. Otherwise, the content of this object is  
 29626 undefined.

29627 **ERRORS**

29628 These functions may fail if:

29629 [EINVAL] The value specified by at least one of the arguments is invalid.

29630 **EXAMPLES**

29631 None.

29632 **APPLICATION USAGE**

29633 None.

29634 **RATIONALE**

29635 None.

29636 **FUTURE DIRECTIONS**

29637 None.

29638 **SEE ALSO**

29639 *fork()*, *posix\_trace\_attr\_init()*, *posix\_trace\_create()*, *posix\_trace\_flush()*, *posix\_trace\_get\_attr()*,  
29640 *posix\_trace\_getnext\_event()*, *posix\_trace\_start()*, **posix\_trace\_status\_info Structure**,  
29641 *posix\_trace\_timedgetnext\_event()*, the Base Definitions volume of IEEE Std 1003.1-200x, <**trace.h**>

29642 **CHANGE HISTORY**

29643 First released in Issue 6. Derived from IEEE Std 1003.1q-2000.

29644 **NAME**

29645        posix\_trace\_attr\_getlogfullpolicy — retrieve and set the behavior of a trace stream (**TRACING**)

29646 **SYNOPSIS**

29647 TRC     #include <trace.h>

29648 TRC TRL int posix\_trace\_attr\_getlogfullpolicy(const trace\_attr\_t \*restrict attr,  
29649           int \*restrict logpolicy);

29650

29651 **DESCRIPTION**

29652        Refer to *posix\_trace\_attr\_getinherited()*.



## 29653 NAME

29654 posix\_trace\_attr\_getlogsize, posix\_trace\_attr\_getmaxdatasize,  
 29655 posix\_trace\_attr\_getmaxsystemeventsize, posix\_trace\_attr\_getmaxusereventsize,  
 29656 posix\_trace\_attr\_getstreamsize, posix\_trace\_attr\_setlogsize, posix\_trace\_attr\_setmaxdatasize,  
 29657 posix\_trace\_attr\_setstreamsize — retrieve and set trace stream size attributes (TRACING)

## 29658 SYNOPSIS

```
29659 TRC #include <sys/types.h>
29660 #include <trace.h>

29661 TRC TRL int posix_trace_attr_getlogsize(const trace_attr_t *restrict attr,
29662 size_t *restrict logsize);
29663 TRC int posix_trace_attr_getmaxdatasize(const trace_attr_t *restrict attr,
29664 size_t *restrict maxdatasize);
29665 int posix_trace_attr_getmaxsystemeventsize(
29666 const trace_attr_t *restrict attr,
29667 size_t *restrict eventsize);
29668 int posix_trace_attr_getmaxusereventsize(
29669 const trace_attr_t *restrict attr,
29670 size_t data_len, size_t *restrict eventsize);
29671 int posix_trace_attr_getstreamsize(const trace_attr_t *restrict attr,
29672 size_t *restrict streamsize);
29673 TRC TRL int posix_trace_attr_setlogsize(trace_attr_t *attr,
29674 size_t logsize);
29675 TRC int posix_trace_attr_setmaxdatasize(trace_attr_t *attr,
29676 size_t maxdatasize);
29677 int posix_trace_attr_setstreamsize(trace_attr_t *attr,
29678 size_t streamsize);
29679
```

## 29680 DESCRIPTION

29681 TRL The *posix\_trace\_attr\_getlogsize()* function shall copy the log size, in bytes, from the *log-max-size*  
 29682 attribute of the attributes object pointed to by the *attr* argument into the variable pointed to by  
 29683 the *logsize* argument. This log size is the maximum total of bytes that shall be allocated for  
 29684 system and user trace events in the trace log. The default value for the *log-max-size* attribute is  
 29685 implementation-defined.

29686 The *posix\_trace\_attr\_setlogsize()* function shall set the maximum allowed size, in bytes, in the  
 29687 *log-max-size* attribute of the attributes object pointed to by the *attr* argument, using the size value  
 29688 supplied by the *logsize* argument.

29689 The trace log size shall be used if the *log-full-policy* attribute is set to `POSIX_TRACE_LOOP` or  
 29690 `POSIX_TRACE_UNTIL_FULL`. If the *log-full-policy* attribute is set to `POSIX_TRACE_APPEND`,  
 29691 the implementation shall ignore the *log-max-size* attribute.

29692 The *posix\_trace\_attr\_getmaxdatasize()* function shall copy the maximum user trace event data  
 29693 size, in bytes, from the *max-data-size* attribute of the attributes object pointed to by the *attr*  
 29694 argument into the variable pointed to by the *maxdatasize* argument. The default value for the  
 29695 *max-data-size* attribute is implementation-defined.

29696 The *posix\_trace\_attr\_getmaxsystemeventsize()* function shall calculate the maximum memory size,  
 29697 in bytes, required to store a single system trace event. This value is calculated for the trace  
 29698 stream attributes object pointed to by the *attr* argument and is returned in the variable pointed  
 29699 to by the *eventsize* argument.

29700 The values returned as the maximum memory sizes of the user and system trace events shall be  
 29701 such that if the sum of the maximum memory sizes of a set of the trace events that may be  
 29702 recorded in a trace stream is less than or equal to the *stream-min-size* attribute of that trace  
 29703 stream, the system provides the necessary resources for recording all those trace events, without  
 29704 loss.

29705 The *posix\_trace\_attr\_getmaxusereventsize()* function shall calculate the maximum memory size, in  
 29706 bytes, required to store a single user trace event generated by a call to *posix\_trace\_event()* with a  
 29707 *data\_len* parameter equal to the *data\_len* value specified in this call. This value is calculated for  
 29708 the trace stream attributes object pointed to by the *attr* argument and is returned in the variable  
 29709 pointed to by the *eventsize* argument.

29710 The *posix\_trace\_attr\_getstreamsize()* function shall copy the stream size, in bytes, from the  
 29711 *stream-min-size* attribute of the attributes object pointed to by the *attr* argument into the variable  
 29712 pointed to by the *streamsize* argument.

29713 This stream size is the current total memory size reserved for system and user trace events in the  
 29714 trace stream. The default value for the *stream-min-size* attribute is implementation-defined. The  
 29715 stream size refers to memory used to store trace event records. Other stream data (for example,  
 29716 trace attribute values) shall not be included in this size.

29717 The *posix\_trace\_attr\_setmaxdatasize()* function shall set the maximum allowed size, in bytes, in  
 29718 the *max-data-size* attribute of the attributes object pointed to by the *attr* argument, using the size  
 29719 value supplied by the *maxdatasize* argument. This maximum size is the maximum allowed size  
 29720 for the user data argument which may be passed to *posix\_trace\_event()*. The implementation  
 29721 shall be allowed to truncate data passed to *trace\_user\_event* which is longer than *maxdatasize*.

29722 The *posix\_trace\_attr\_setstreamsize()* function shall set the minimum allowed size, in bytes, in the  
 29723 *stream-min-size* attribute of the attributes object pointed to by the *attr* argument, using the size  
 29724 value supplied by the *streamsize* argument.

29725 **RETURN VALUE**

29726 Upon successful completion, these functions shall return a value of zero. Otherwise, they shall  
 29727 return the corresponding error number.

29728 TRL The *posix\_trace\_attr\_getlogsize()* function stores the maximum trace log allowed size in the object  
 29729 pointed to by *logsize*, if successful.

29730 The *posix\_trace\_attr\_getmaxdatasize()* function stores the maximum trace event record memory  
 29731 size in the object pointed to by *maxdatasize*, if successful.

29732 The *posix\_trace\_attr\_getmaxsystemeventsize()* function stores the maximum memory size to store  
 29733 a single system trace event in the object pointed to by *eventsize*, if successful.

29734 The *posix\_trace\_attr\_getmaxusereventsize()* function stores the maximum memory size to store a  
 29735 single user trace event in the object pointed to by *eventsize*, if successful.

29736 The *posix\_trace\_attr\_getstreamsize()* function stores the maximum trace stream allowed size in  
 29737 the object pointed to by *streamsize*, if successful.

29738 **ERRORS**

29739 These functions may fail if:

29740 [EINVAL] The value specified by one of the arguments is invalid.

29741 **EXAMPLES**

29742 None.

29743 **APPLICATION USAGE**

29744 None.

29745 **RATIONALE**

29746 None.

29747 **FUTURE DIRECTIONS**

29748 None.

29749 **SEE ALSO**29750 *posix\_trace\_attr\_init()*, *posix\_trace\_create()*, *posix\_trace\_event()*, *posix\_trace\_get\_attr()*, the Base

29751 Definitions volume of IEEE Std 1003.1-200x, &lt;sys/types.h&gt;, &lt;trace.h&gt;

29752 **CHANGE HISTORY**

29753 First released in Issue 6. Derived from the IEEE Std 1003.1q-2000.

29754 **NAME**

29755 posix\_trace\_attr\_getmaxdatasize, posix\_trace\_attr\_getmaxsystemeventszize,  
29756 posix\_trace\_attr\_getmaxusereventszize — retrieve and set trace stream size attributes  
29757 **(TRACING)**

29758 **SYNOPSIS**

```
29759 TRC #include <sys/types.h>  
29760 #include <trace.h>  
  
29761 TRC int posix_trace_attr_getmaxdatasize(const trace_attr_t *restrict attr,  
29762 size_t *restrict maxdatasize);  
29763 int posix_trace_attr_getmaxsystemeventszize(  
29764 const trace_attr_t *restrict attr,  
29765 size_t *restrict eventszize);  
29766 int posix_trace_attr_getmaxusereventszize(  
29767 const trace_attr_t *restrict attr,  
29768 size_t data_len, size_t *restrict eventszize);  
29769
```

29770 **DESCRIPTION**

29771 Refer to *posix\_trace\_attr\_getlogszize()*.

29772 **NAME**

29773 `posix_trace_attr_getname` — retrieve and set information about a trace stream (**TRACING**)

29774 **SYNOPSIS**

29775 TRC `#include <trace.h>`

```
29776 int posix_trace_attr_getname(const trace_attr_t *attr,  
29777 char *tracename);
```

29778

29779 **DESCRIPTION**

29780 Refer to `posix_trace_attr_getclockres()`.

29781 **NAME**

29782        posix\_trace\_attr\_getstreamfullpolicy — retrieve and set the behavior of a trace stream  
29783        **(TRACING)**

29784 **SYNOPSIS**

```
29785 TRC     #include <trace.h>
29786           int posix_trace_attr_getstreamfullpolicy(const trace_attr_t *attr,
29787                                                    int *streampolicy);
29788
```

29789 **DESCRIPTION**

29790        Refer to *posix\_trace\_attr\_getinherited()*.

29791 **NAME**29792 `posix_trace_attr_getstreamsize` — retrieve and set trace stream size attributes (**TRACING**)29793 **SYNOPSIS**29794 TRC `#include <sys/types.h>`29795 `#include <trace.h>`29796 `int posix_trace_attr_getstreamsize(const trace_attr_t *restrict attr,`  
29797 `size_t *restrict streamsize);`

29798

29799 **DESCRIPTION**29800 Refer to `posix_trace_attr_getlogsize()`.

29801 **NAME**

29802        posix\_trace\_attr\_init — trace stream attributes object initialization (**TRACING**)

29803 **SYNOPSIS**

29804 TRC     #include <trace.h>

29805        int posix\_trace\_attr\_init(trace\_attr\_t \*attr);

29806

29807 **DESCRIPTION**

29808        Refer to *posix\_trace\_attr\_destroy()*.



29809 **NAME**

29810 `posix_trace_attr_setinherited` — retrieve and set the behavior of a trace stream (**TRACING**)

29811 **SYNOPSIS**

29812 TRC `#include <trace.h>`

29813 TRC TRI `int posix_trace_attr_setinherited(trace_attr_t *attr,`  
29814 `int inheritancepolicy);`

29815

29816 **DESCRIPTION**

29817 Refer to `posix_trace_attr_getinherited()`.

29818 **NAME**

29819        posix\_trace\_attr\_setlogfullpolicy — retrieve and set the behavior of a trace stream (**TRACING**)

29820 **SYNOPSIS**

29821 TRC     #include <trace.h>

29822 TRC TRL int posix\_trace\_attr\_setlogfullpolicy(trace\_attr\_t \*attr,  
29823           int logpolicy);

29824

29825 **DESCRIPTION**

29826        Refer to *posix\_trace\_attr\_getinherited()*.

29827 **NAME**29828        posix\_trace\_attr\_setlogsize — retrieve and set trace stream size attributes (**TRACING**)29829 **SYNOPSIS**

29830 TRC     #include &lt;sys/types.h&gt;

29831        #include &lt;trace.h&gt;

29832 TRC TRL int posix\_trace\_attr\_setlogsize(trace\_attr\_t \*attr,  
29833        size\_t logsize);

29834

29835 **DESCRIPTION**29836        Refer to *posix\_trace\_attr\_getlogsize()*.

29837 **NAME**

29838        posix\_trace\_attr\_setmaxdatasize — retrieve and set trace stream size attributes (**TRACING**)

29839 **SYNOPSIS**

29840 TRC     #include <sys/types.h>

29841        #include <trace.h>

```
29842        int posix_trace_attr_setmaxdatasize(trace_attr_t *attr,  
29843                                            size_t maxdatasize);
```

29844

29845 **DESCRIPTION**

29846        Refer to *posix\_trace\_attr\_getlogsize()*.

29847 **NAME**

29848        posix\_trace\_attr\_setname — retrieve and set information about a trace stream (**TRACING**)

29849 **SYNOPSIS**

29850 TRC     #include <trace.h>

```
29851        int posix_trace_attr_setname(trace_attr_t *attr,  
29852            const char *tracename);
```

29853

29854 **DESCRIPTION**

29855        Refer to *posix\_trace\_attr\_getclockres()*.

29856 **NAME**

29857        posix\_trace\_attr\_setstreamfullpolicy — retrieve and set the behavior of a trace stream  
29858        **(TRACING)**

29859 **SYNOPSIS**

```
29860 TRC     #include <trace.h>
```

```
29861 TRC TRL int posix_trace_attr_setlogfullpolicy(trace_attr_t *attr,  
29862                                           int logpolicy);
```

29863

29864 **DESCRIPTION**

29865        Refer to *posix\_trace\_attr\_getinherited()*.

29866 **NAME**

29867        posix\_trace\_attr\_setstreamsize — retrieve and set trace stream size attributes (**TRACING**)

29868 **SYNOPSIS**

29869 TRC     #include <sys/types.h>

29870        #include <trace.h>

29871        int posix\_trace\_attr\_setstreamsize(trace\_attr\_t \*attr,  
29872                                    size\_t streamsize);

29873

29874 **DESCRIPTION**

29875        Refer to *posix\_trace\_attr\_getlogsize()*.

29876 **NAME**

29877 posix\_trace\_clear — clear trace stream and trace log (TRACING)

29878 **SYNOPSIS**

29879 TRC #include <sys/types.h>

29880 #include <trace.h>

29881 int posix\_trace\_clear(trace\_id\_t trid);

29882

29883 **DESCRIPTION**

29884 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 shall be reused, the mapping of trace event type identifiers to trace event names shall be unchanged, and the trace stream status shall remain unchanged (that is, if it was running, it remains running and if it was suspended, it remains suspended).

29889 All trace events in the trace stream recorded before the call to *posix\_trace\_clear()* shall be 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.

29893 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.

29906 **RETURN VALUE**

29907 Upon successful completion, the *posix\_trace\_clear()* function shall return a value of zero. 29908 Otherwise, it shall return the corresponding error number.

29909 **ERRORS**

29910 The *posix\_trace\_clear()* function shall fail if:

29911 [EINVAL] The value of the *trid* argument does not correspond to an active trace stream.

29912 **EXAMPLES**

29913 None.

29914 **APPLICATION USAGE**

29915 None.

29916 **RATIONALE**

29917 None.

29918 **FUTURE DIRECTIONS**

29919 None.



29920 **SEE ALSO**

29921 *posix\_trace\_attr\_init()*, *posix\_trace\_create()*, *posix\_trace\_flush()*, *posix\_trace\_get\_attr()*, the Base  
29922 Definitions volume of IEEE Std 1003.1-200x, <sys/types.h>, <trace.h>

29923 **CHANGE HISTORY**

29924 First released in Issue 6. Derived from IEEE Std 1003.1q-2000. |

29925 IEEE PASC Interpretation 1003.1 #123 is applied. |

29926 **NAME**

29927 posix\_trace\_close, posix\_trace\_open, posix\_trace\_rewind — trace log management (**TRACING**)

29928 **SYNOPSIS**

29929 TRC TRL #include <trace.h>

```
29930 int posix_trace_close(trace_id_t trid);
29931 int posix_trace_open(int file_desc, trace_id_t *trid);
29932 int posix_trace_rewind(trace_id_t trid);
29933
```

29934 **DESCRIPTION**

29935 The *posix\_trace\_close()* function shall deallocate the trace log identifier indicated by *trid*, and all  
 29936 of its associated resources. If there is no valid trace log pointed to by the *trid*, this function shall  
 29937 fail.

29938 The *posix\_trace\_open()* function shall allocate the necessary resources and establishes the  
 29939 connection between a trace log identified by the *file\_desc* argument and a trace stream identifier  
 29940 identified by the object pointed to by the *trid* argument. The *file\_desc* argument should be a valid  
 29941 open file descriptor that corresponds to a trace log. The *file\_desc* argument shall be open for  
 29942 reading. The current trace event timestamp, which specifies the timestamp of the trace event  
 29943 that will be read by the next call to *posix\_trace\_getnext\_event()*, shall be set to the timestamp of  
 29944 the oldest trace event recorded in the trace log identified by *trid*.

29945 The *posix\_trace\_open()* function shall return a trace stream identifier in the variable pointed to by  
 29946 the *trid* argument, that may only be used by the following functions:

29947	<i>posix_trace_close()</i>	<i>posix_trace_get_attr()</i>
29948	<i>posix_trace_eventid_equal()</i>	<i>posix_trace_get_status()</i>
29949	<i>posix_trace_eventid_get_name()</i>	<i>posix_trace_getnext_event()</i>
29950	<i>posix_trace_eventtypelist_getnext_id()</i>	<i>posix_trace_rewind()</i>
29951	<i>posix_trace_eventtypelist_rewind()</i>	

29952 In particular, notice that the operations normally used by a trace controller process, such as  
 29953 *posix\_trace\_start()*, *posix\_trace\_stop()*, or *posix\_trace\_shutdown()*, cannot be invoked using the  
 29954 trace stream identifier returned by the *posix\_trace\_open()* function.

29955 The *posix\_trace\_rewind()* function shall reset the current trace event timestamp, which specifies  
 29956 the timestamp of the trace event that will be read by the next call to *posix\_trace\_getnext\_event()*,  
 29957 to the timestamp of the oldest trace event recorded in the trace log identified by *trid*.

29958 **RETURN VALUE**

29959 Upon successful completion, these functions shall return a value of zero. Otherwise, they shall  
 29960 return the corresponding error number.

29961 If successful, the *posix\_trace\_open()* function stores the trace stream identifier value in the object  
 29962 pointed to by *trid*.

29963 **ERRORS**

29964 The *posix\_trace\_open()* function shall fail if:

- 29965 [EINTR] The operation was interrupted by a signal and thus no trace log was opened.
- 29966 [EINVAL] The object pointed to by *file\_desc* does not correspond to a valid trace log.

29967 The *posix\_trace\_close()* and *posix\_trace\_rewind()* functions may fail if:

- 29968 [EINVAL] The object pointed to by *trid* does not correspond to a valid trace log.

29969 **EXAMPLES**

29970 None.

29971 **APPLICATION USAGE**

29972 None.

29973 **RATIONALE**

29974 None.

29975 **FUTURE DIRECTIONS**

29976 None.

29977 **SEE ALSO**29978 *posix\_trace\_get\_attr()*, *posix\_trace\_get\_filter()*, *posix\_trace\_getnext\_event()*, the Base Definitions29979 volume of IEEE Std 1003.1-200x, <**trace.h**>29980 **CHANGE HISTORY**

29981 First released in Issue 6. Derived from IEEE Std 1003.1q-2000. |

29982 IEEE PASC Interpretation 1003.1 #123 is applied. |

29983 **NAME**

29984 posix\_trace\_create, posix\_trace\_create\_withlog, posix\_trace\_flush, posix\_trace\_shutdown —  
 29985 trace stream initialization, flush, and shutdown from a process (**TRACING**)

29986 **SYNOPSIS**

```

29987 TRC #include <sys/types.h>
29988 #include <trace.h>

29989 int posix_trace_create(pid_t pid,
29990 const trace_attr_t *restrict attr,
29991 trace_id_t *restrict trid);
29992 TRC TRL int posix_trace_create_withlog(pid_t pid,
29993 const trace_attr_t *restrict attr, int file_desc,
29994 trace_id_t *restrict trid);
29995 int posix_trace_flush(trace_id_t trid);
29996 TRC int posix_trace_shutdown(trace_id_t trid);
29997
    
```

29998 **DESCRIPTION**

29999 The *posix\_trace\_create()* function shall create an active trace stream. It allocates all the resources  
 30000 needed by the trace stream being created for tracing the process specified by *pid* in accordance  
 30001 with the *attr* argument. The *attr* argument represents the initial attributes of the trace stream and  
 30002 shall have been initialized by the function *posix\_trace\_attr\_init()* prior the *posix\_trace\_create()*  
 30003 call. If the argument *attr* is NULL, the default attributes shall be used. The *attr* attributes object  
 30004 shall be manipulated through a set of functions described in the *posix\_trace\_attr* family of  
 30005 functions. If the attributes of the object pointed to by *attr* are modified later, the attributes of the  
 30006 trace stream shall not be affected. The *creation-time* attribute of the newly created trace stream  
 30007 shall be set to the value of the system clock, if the Timers option is not supported, or to the value  
 30008 of the CLOCK\_REALTIME clock, if the Timers option is supported.

30009 The *pid* argument represents the target process to be traced. If the process executing this  
 30010 function does not have appropriate privileges to trace the process identified by *pid*, an error shall  
 30011 be returned. If the *pid* argument is zero, the calling process shall be traced.

30012 The *posix\_trace\_create()* function shall store the trace stream identifier of the new trace stream in  
 30013 the object pointed to by the *trid* argument. This trace stream identifier shall be used in  
 30014 subsequent calls to control tracing. The *trid* argument may only be used by the following  
 30015 functions:

30016	<i>posix_trace_clear()</i>	<i>posix_trace_getnext_event()</i>
30017	<i>posix_trace_eventid_equal()</i>	<i>posix_trace_shutdown()</i>
30018	<i>posix_trace_eventid_get_name()</i>	<i>posix_trace_start()</i>
30019	<i>posix_trace_eventtypelist_getnext_id()</i>	<i>posix_trace_stop()</i>
30020	<i>posix_trace_eventtypelist_rewind()</i>	<i>posix_trace_timedgetnext_event()</i>
30021	<i>posix_trace_get_attr()</i>	<i>posix_trace_trid_eventid_open()</i>
30022	<i>posix_trace_get_status()</i>	<i>posix_trace_trygetnext_event()</i>

30023 TEF If the Trace Event Filter option is supported, the following additional functions may use the *trid*  
 30024 argument:

```

30025 posix_trace_get_filter() posix_trace_set_filter()
    
```

30026

30027 In particular, notice that the operations normally used by a trace analyzer process, such as  
 30028 *posix\_trace\_rewind()* or *posix\_trace\_close()*, cannot be invoked using the trace stream identifier  
 30029 returned by the *posix\_trace\_create()* function.

30030 TEF A trace stream shall be created in a suspended state. If the Trace Event Filter option is  
 30031 supported, its trace event type filter shall be empty.

30032 The *posix\_trace\_create()* function may be called multiple times from the same or different  
 30033 processes, with the system-wide limit indicated by the runtime invariant value  
 30034 {TRACE\_SYS\_MAX}, which has the minimum value {\_POSIX\_TRACE\_SYS\_MAX}.

30035 The trace stream identifier returned by the *posix\_trace\_create()* function in the argument pointed  
 30036 to by *trid* is valid only in the process that made the function call. If it is used from another  
 30037 process, that is a child process, in functions defined in IEEE Std 1003.1-200x, these functions shall  
 30038 return with the error [EINVAL].

30039 TRL The *posix\_trace\_create\_withlog()* function shall be equivalent to *posix\_trace\_create()*, except that it |  
 30040 associates a trace log with this stream. The *file\_desc* argument shall be the file descriptor |  
 30041 designating the trace log destination. The function shall fail if this file descriptor refers to a file |  
 30042 with a file type that is not compatible with the log policy associated with the trace log. The list of |  
 30043 the appropriate file types that are compatible with each log policy is implementation-defined. |

30044 The *posix\_trace\_create\_withlog()* function shall return in the parameter pointed to by *trid* the trace  
 30045 stream identifier, which uniquely identifies the newly created trace stream, and shall be used in  
 30046 subsequent calls to control tracing. The *trid* argument may only be used by the following  
 30047 functions:

30048	<i>posix_trace_clear()</i>	<i>posix_trace_getnext_event()</i>
30049	<i>posix_trace_eventid_equal()</i>	<i>posix_trace_shutdown()</i>
30050	<i>posix_trace_eventid_get_name()</i>	<i>posix_trace_start()</i>
30051	<i>posix_trace_eventtypelist_getnext_id()</i>	<i>posix_trace_stop()</i>
30052	<i>posix_trace_eventtypelist_rewind()</i>	<i>posix_trace_timedgetnext_event()</i>
30053	<i>posix_trace_flush()</i>	<i>posix_trace_trid_eventid_open()</i>
30054	<i>posix_trace_get_attr()</i>	<i>posix_trace_trygetnext_event()</i>
30055	<i>posix_trace_get_status()</i>	

30056

30057 TRL TEF If the Trace Event Filter option is supported, the following additional functions may use the *trid*  
 30058 argument:

30059 *posix\_trace\_get\_filter()*    *posix\_trace\_set\_filter()*

30060

30061 TRL In particular, notice that the operations normally used by a trace analyzer process, such as  
 30062 *posix\_trace\_rewind()* or *posix\_trace\_close()*, cannot be invoked using the trace stream identifier  
 30063 returned by the *posix\_trace\_create\_withlog()* function.

30064 The *posix\_trace\_flush()* function shall initiate a flush operation which copies the contents of the |  
 30065 trace stream identified by the argument *trid* into the trace log associated with the trace stream at |  
 30066 the creation time. If no trace log has been associated with the trace stream pointed to by *trid*, this |  
 30067 function shall return an error. The termination of the flush operation can be polled by the |  
 30068 *posix\_trace\_get\_status()* function. During the flush operation, it shall be possible to trace new |  
 30069 trace events up to the point when the trace stream becomes full. After flushing is completed, the |  
 30070 space used by the flushed trace events shall be available for tracing new trace events.

30071 If flushing the trace stream causes the resulting trace log to become full, the trace log full policy  
 30072 shall be applied. If the trace *log-full-policy* attribute is set, the following occurs:

30073 POSIX\_TRACE\_UNTIL\_FULL  
 30074 The trace events that have not yet been flushed shall be discarded. |

30075 POSIX\_TRACE\_LOOP  
 30076 The trace events that have not yet been flushed shall be written to the beginning of the trace |  
 30077 log, overwriting previous trace events stored there.

30078 POSIX\_TRACE\_APPEND  
 30079 The trace events that had not yet been flushed shall be appended to the trace log.  
 30080

30081 The *posix\_trace\_shutdown()* function shall stop the tracing of trace events in the trace stream  
 30082 identified by *trid*, as if *posix\_trace\_stop()* had been invoked. The *posix\_trace\_shutdown()* function  
 30083 shall free all the resources associated with the trace stream.

30084 The *posix\_trace\_shutdown()* function shall not return until all the resources associated with the  
 30085 trace stream have been freed. When the *posix\_trace\_shutdown()* function returns, the *trid*  
 30086 argument becomes an invalid trace stream identifier. A call to this function shall unconditionally  
 30087 deallocate the resources regardless of whether all trace events have been retrieved by the  
 30088 analyzer process. Any thread blocked on one of the *trace\_getnext\_event()* functions (which  
 30089 specified this *trid*) before this call is unblocked with the error [EINVAL].

30090 If the process exits, invokes an *exec()* call, or is terminated, the trace streams that the process had  
 30091 created and that have not yet been shut down, shall be automatically shut down as if an explicit  
 30092 call were made to the *posix\_trace\_shutdown()* function.

30093 TRL For an active trace stream with log, when the *posix\_trace\_shutdown()* function is called, all trace  
 30094 events that have not yet been flushed to the trace log shall be flushed, as in the  
 30095 *posix\_trace\_flush()* function, and the trace log shall be closed.

30096 When a trace log is closed, all the information that may be retrieved later from the trace log |  
 30097 through the trace interface shall have been written to the trace log. This information includes the |  
 30098 trace attributes, the list of trace event types (with the mapping between trace event names and |  
 30099 trace event type identifiers), and the trace status.

30100 In addition, unspecified information shall be written to the trace log to allow detection of a valid |  
 30101 trace log during the *posix\_trace\_open()* operation. |

30102 The *posix\_trace\_shutdown()* function shall not return until all trace events have been flushed.

30103 **RETURN VALUE**

30104 Upon successful completion, these functions shall return a value of zero. Otherwise, they shall  
 30105 return the corresponding error number.

30106 TRL The *posix\_trace\_create()* and *posix\_trace\_create\_withlog()* functions store the trace stream identifier  
 30107 value in the object pointed to by *trid*, if successful.

30108 **ERRORS**

30109 TRL The *posix\_trace\_create()* and *posix\_trace\_create\_withlog()* functions shall fail if:

30110 [EAGAIN] No more trace streams can be started now. {TRACE\_SYS\_MAX} has been  
 30111 exceeded.

30112 [EINTR] The operation was interrupted by a signal. No trace stream was created.

30113 [EINVAL] One or more of the trace parameters specified by the *attr* parameter is invalid.

30114	[ENOMEM]	The implementation does not currently have sufficient memory to create the trace stream with the specified parameters.
30115		
30116	[EPERM]	The caller does not have appropriate privilege to trace the process specified by <i>pid</i> .
30117		
30118	[ESRCH]	The <i>pid</i> argument does not refer to an existing process.
30119	TRL	The <i>posix_trace_create_withlog()</i> function shall fail if:
30120	[EBADF]	The <i>file_desc</i> argument is not a valid file descriptor open for writing.
30121	[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.
30122		
30123	[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.
30124		
30125		
30126	TRL	The <i>posix_trace_flush()</i> and <i>posix_trace_shutdown()</i> functions shall fail if:
30127	[EINVAL]	The value of the <i>trid</i> argument does not correspond to an active trace stream with log.
30128		
30129	[EFBIG]	The trace log file has attempted to exceed an implementation-defined maximum file size.
30130		
30131	[ENOSPC]	No space left on device.
30132		

**30133 EXAMPLES**

30134 None.

**30135 APPLICATION USAGE**

30136 None.

**30137 RATIONALE**

30138 None.

**30139 FUTURE DIRECTIONS**

30140 None.

**30141 SEE ALSO**

30142 *clock\_getres()*, *exec*, *posix\_trace\_attr\_init()*, *posix\_trace\_clear()*, *posix\_trace\_close()*,  
 30143 *posix\_trace\_eventid\_equal()*, *posix\_trace\_eventtypelist\_getnext\_id()*, *posix\_trace\_flush()*,  
 30144 *posix\_trace\_get\_attr()*, *posix\_trace\_get\_filter()*, *posix\_trace\_get\_status()*, *posix\_trace\_getnext\_event()*,  
 30145 *posix\_trace\_open()*, *posix\_trace\_rewind()*, *posix\_trace\_set\_filter()*, *posix\_trace\_shutdown()*,  
 30146 *posix\_trace\_start()*, *posix\_trace\_timedgetnext\_event()*, *posix\_trace\_trid\_eventid\_open()*,  
 30147 *posix\_trace\_start()*, *time()*, the Base Definitions volume of IEEE Std 1003.1-200x, <sys/types.h>,  
 30148 <trace.h>

**30149 CHANGE HISTORY**

30150 First released in Issue 6. Derived from IEEE Std 1003.1q-2000.

30151 NAME

30152 posix\_trace\_event, posix\_trace\_eventid\_open — trace functions for instrumenting application  
 30153 code (TRACING)

30154 SYNOPSIS

```
30155 TRC #include <sys/types.h>
30156 #include <trace.h>

30157 void posix_trace_event(trace_event_id_t event_id,
30158     const void *restrictdata_ptr, size_t data_len);
30159 int posix_trace_eventid_open(const char *restrict event_name,
30160     trace_event_id_t *restrict event_id);
30161
```

30162 DESCRIPTION

30163 The *posix\_trace\_event()* function shall record the *event\_id* and the user data pointed to by *data\_ptr*  
 30164 in the trace stream into which the calling process is being traced and in which *event\_id* is not  
 30165 filtered out. If the total size of the user trace event data represented by *data\_len* is not greater  
 30166 than the declared maximum size for user trace event data, then the *truncation-status* attribute of  
 30167 the trace event recorded is POSIX\_TRACE\_NOT\_TRUNCATED. Otherwise, the user trace event  
 30168 data is truncated to this declared maximum size and the *truncation-status* attribute of the trace  
 30169 event recorded is POSIX\_TRACE\_TRUNCATED\_RECORD.

30170 If there is no trace stream created for the process or if the created trace stream is not running or if  
 30171 the trace event specified by *event\_id* is filtered out in the trace stream, the *posix\_trace\_event()*  
 30172 function shall have no effect.

30173 The *posix\_trace\_eventid\_open()* function shall associate a user trace event name with a trace event  
 30174 type identifier for the calling process. The trace event name is the string pointed to by the  
 30175 argument *event\_name*. It shall have a maximum of {TRACE\_EVENT\_NAME\_MAX} characters  
 30176 (which has the minimum value {POSIX\_TRACE\_EVENT\_NAME\_MAX}). The number of user  
 30177 trace event type identifiers that can be defined for any given process is limited by the maximum  
 30178 value {TRACE\_USER\_EVENT\_MAX}, which has the minimum value  
 30179 {POSIX\_TRACE\_USER\_EVENT\_MAX}.

30180 If the Trace Inherit option is not supported, the *posix\_trace\_eventid\_open()* function shall  
 30181 associate the user trace event name pointed to by the *event\_name* argument with a trace event  
 30182 type identifier that is unique for the traced process, and is returned in the variable pointed to by  
 30183 the *event\_id* argument. If the user trace event name has already been mapped for the traced  
 30184 process, then the previously assigned trace event type identifier shall be returned. If the per-  
 30185 process user trace event name limit represented by {TRACE\_USER\_EVENT\_MAX} has been  
 30186 reached, the pre-defined POSIX\_TRACE\_UNNAMED\_USEREVENT (see Table 2-7 (on page  
 30187 529)) user trace event shall be returned.

30188 TRI If the Trace Inherit option is supported, the *posix\_trace\_eventid\_open()* function shall associate the  
 30189 user trace event name pointed to by the *event\_name* argument with a trace event type identifier  
 30190 that is unique for all the processes being traced in this same trace stream, and is returned in the  
 30191 variable pointed to by the *event\_id* argument. If the user trace event name has already been  
 30192 mapped for the traced processes, then the previously assigned trace event type identifier shall be  
 30193 returned. If the per-process user trace event name limit represented by  
 30194 {TRACE\_USER\_EVENT\_MAX} has been reached, the pre-defined  
 30195 POSIX\_TRACE\_UNNAMED\_USEREVENT (Table 2-7 (on page 529)) user trace event shall be  
 30196 returned.

30197 **Note:** The above procedure, together with the fact that multiple processes can only be traced into the  
 30198 same trace stream by inheritance, ensure that all the processes that are traced into a trace  
 30199 stream have the same mapping of trace event names to trace event type identifiers.



30200

30201 If there is no trace stream created, the *posix\_trace\_eventid\_open()* function shall store this  
30202 information for future trace streams created for this process.

**30203 RETURN VALUE**

30204 No return value is defined for the *posix\_trace\_event()* function.

30205 Upon successful completion, the *posix\_trace\_eventid\_open()* function shall return a value of zero.  
30206 Otherwise, it shall return the corresponding error number. The *posix\_trace\_eventid\_open()*  
30207 function stores the trace event type identifier value in the object pointed to by *event\_id*, if  
30208 successful.

**30209 ERRORS**

30210 The *posix\_trace\_eventid\_open()* function shall fail if:

30211 [ENAMETOOLONG]

30212 The size of the name pointed to by *event\_name* argument was longer than the  
30213 implementation-defined value {TRACE\_EVENT\_NAME\_MAX}.

**30214 EXAMPLES**

30215 None.

**30216 APPLICATION USAGE**

30217 None.

**30218 RATIONALE**

30219 None.

**30220 FUTURE DIRECTIONS**

30221 None.

**30222 SEE ALSO**

30223 *posix\_trace\_start()*, *posix\_trace\_trid\_eventid\_open()*, the Base Definitions volume of  
30224 IEEE Std 1003.1-200x, <sys/types.h>, <trace.h>

**30225 CHANGE HISTORY**

30226 First released in Issue 6. Derived from IEEE Std 1003.1q-2000.

30227 IEEE PASC Interpretation 1003.1 #123 is applied.

30228 IEEE PASC Interpretation 1003.1 #127 is applied, correcting some editorial errors in the names of  
30229 the *posix\_trace\_eventid\_open()* function and the *event\_id* argument.

30230 NAME

30231 posix\_trace\_eventid\_equal, posix\_trace\_eventid\_get\_name, posix\_trace\_trid\_eventid\_open —  
 30232 manipulate trace event type identifier (TRACING)

30233 SYNOPSIS

```
30234 TRC #include <trace.h>
30235 int posix_trace_eventid_equal(trace_id_t trid, trace_event_id_t event1, |
30236 trace_event_id_t event2); |
30237 int posix_trace_eventid_get_name(trace_id_t trid, |
30238 trace_event_id_t event, char *event_name); |
30239 TRC TEF int posix_trace_trid_eventid_open(trace_id_t trid, |
30240 const char *restrict event_name, |
30241 trace_event_id_t *restrict event); |
30242
```

30243 DESCRIPTION

30244 The *posix\_trace\_eventid\_equal()* function shall compare the trace event type identifiers *event1* and  
 30245 *event2* from the same trace stream or the same trace log identified by the *trid* argument. If the  
 30246 trace event type identifiers *event1* and *event2* are from different trace streams, the return value  
 30247 shall be unspecified.

30248 The *posix\_trace\_eventid\_get\_name()* function shall return in the argument pointed to by  
 30249 *event\_name*, the trace event name associated with the trace event type identifier identified by the  
 30250 argument *event*, for the trace stream or for the trace log identified by the *trid* argument. The  
 30251 name of the trace event shall have a maximum of {TRACE\_EVENT\_NAME\_MAX} characters  
 30252 (which has the minimum value {\_POSIX\_TRACE\_EVENT\_NAME\_MAX}). Successive calls to  
 30253 this function with the same trace event type identifier and the same trace stream identifier shall  
 30254 return the same event name.

30255 TEF The *posix\_trace\_trid\_eventid\_open()* function shall associate a user trace event name with a trace  
 30256 event type identifier for a given trace stream. The trace stream is identified by the *trid* argument,  
 30257 and it shall be an active trace stream. The trace event name is the string pointed to by the  
 30258 argument *event\_name*. It shall have a maximum of {TRACE\_EVENT\_NAME\_MAX} characters  
 30259 (which has the minimum value {\_POSIX\_TRACE\_EVENT\_NAME\_MAX}). The number of user  
 30260 trace event type identifiers that can be defined for any given process is limited by the maximum  
 30261 value {TRACE\_USER\_EVENT\_MAX}, which has the minimum value  
 30262 {\_POSIX\_TRACE\_USER\_EVENT\_MAX}.

30263 If the Trace Inherit option is not supported, the *posix\_trace\_trid\_eventid\_open()* function shall  
 30264 associate the user trace event name pointed to by the *event\_name* argument with a trace event  
 30265 type identifier that is unique for the process being traced in the trace stream identified by the *trid*  
 30266 argument, and is returned in the variable pointed to by the *event* argument. If the user trace  
 30267 event name has already been mapped for the traced process, then the previously assigned trace  
 30268 event type identifier shall be returned. If the per-process user trace event name limit represented  
 30269 by {TRACE\_USER\_EVENT\_MAX} has been reached, the pre-defined  
 30270 POSIX\_TRACE\_UNNAMED\_USEREVENT (see Table 2-7 (on page 529)) user trace event shall  
 30271 be returned.

30272 TEF TRI If the Trace Inherit option is supported, the *posix\_trace\_trid\_eventid\_open()* function shall  
 30273 associate the user trace event name pointed to by the *event\_name* argument with a trace event  
 30274 type identifier that is unique for all the processes being traced in the trace stream identified by  
 30275 the *trid* argument, and is returned in the variable pointed to by the *event* argument. If the user  
 30276 trace event name has already been mapped for the traced processes, then the previously  
 30277 assigned trace event type identifier shall be returned. If the per-process user trace event name  
 30278 limit represented by {TRACE\_USER\_EVENT\_MAX} has been reached, the pre-defined

30279 POSIX\_TRACE\_UNNAMED\_USEREVENT (see Table 2-7 (on page 529)) user trace event shall  
30280 be returned.

### 30281 RETURN VALUE

30282 TEF Upon successful completion, the *posix\_trace\_eventid\_get\_name()* and  
30283 *posix\_trace\_trid\_eventid\_open()* functions shall return a value of zero. Otherwise, they shall return  
30284 the corresponding error number.

30285 The *posix\_trace\_eventid\_equal()* function shall return a non-zero value if *event1* and *event2* are  
30286 equal; otherwise, a value of zero shall be returned. No errors are defined. If either *event1* or  
30287 *event2* are not valid trace event type identifiers for the trace stream specified by *trid* or if the *trid*  
30288 is invalid, the behavior shall be unspecified.

30289 The *posix\_trace\_eventid\_get\_name()* function stores the trace event name value in the object  
30290 pointed to by *event\_name*, if successful.

30291 TEF The *posix\_trace\_trid\_eventid\_open()* function stores the trace event type identifier value in the  
30292 object pointed to by *event*, if successful.

### 30293 ERRORS

30294 TEF The *posix\_trace\_eventid\_get\_name()* and *posix\_trace\_trid\_eventid\_open()* functions shall fail if:

30295 [EINVAL] The *trid* argument was not a valid trace stream identifier.

30296 TEF The *posix\_trace\_trid\_eventid\_open()* function shall fail if:

30297 TEF [ENAMETOOLONG]

30298 The size of the name pointed to by *event\_name* argument was longer than the  
30299 implementation-defined value {TRACE\_EVENT\_NAME\_MAX}.

30300 The *posix\_trace\_eventid\_get\_name()* function shall fail if:

30301 [EINVAL] The trace event type identifier *event* was not associated with any name.

### 30302 EXAMPLES

30303 None.

### 30304 APPLICATION USAGE

30305 None.

### 30306 RATIONALE

30307 None.

### 30308 FUTURE DIRECTIONS

30309 None.

### 30310 SEE ALSO

30311 *posix\_trace\_event()*, *posix\_trace\_getnext\_event()*, the Base Definitions volume of  
30312 IEEE Std 1003.1-200x, <trace.h>

### 30313 CHANGE HISTORY

30314 First released in Issue 6. Derived from IEEE Std 1003.1q-2000.

30315 IEEE PASC Interpretations 1003.1 #123 and #129 are applied.

30316 **NAME**

30317        posix\_trace\_eventid\_get\_name — manipulate trace event type identifier (**TRACING**)

30318 **SYNOPSIS**

30319 TRC     #include <trace.h>

```
30320        int posix_trace_eventid_get_name(trace_id_t trid,  
30321                                    trace_event_id_t event, char *event_name);
```

30322

30323 **DESCRIPTION**

30324        Refer to *posix\_trace\_eventid\_equal()*.

30325 **NAME**30326        posix\_trace\_eventid\_open — trace functions for instrumenting application code (**TRACING**)30327 **SYNOPSIS**

30328 TRC     #include &lt;sys/types.h&gt;

30329        #include &lt;trace.h&gt;

30330        int posix\_trace\_eventid\_open(const char \*restrict event\_name,

30331            trace\_event\_id\_t \*restrict event\_id);

30332

30333 **DESCRIPTION**30334        Refer to *posix\_trace\_event()*.

30335 **NAME**

30336 posix\_trace\_eventset\_add, posix\_trace\_eventset\_del, posix\_trace\_eventset\_empty,  
 30337 posix\_trace\_eventset\_fill, posix\_trace\_eventset\_ismember — manipulate trace event type sets  
 30338 (TRACING)

30339 **SYNOPSIS**

```
30340 TRC TEF #include <trace.h>

30341 int posix_trace_eventset_add(trace_event_id_t event_id,
30342     trace_event_set_t *set);
30343 int posix_trace_eventset_del(trace_event_id_t event_id,
30344     trace_event_set_t *set);
30345 int posix_trace_eventset_empty(trace_event_set_t *set);
30346 int posix_trace_eventset_fill(trace_event_set_t *set, int what);
30347 int posix_trace_eventset_ismember(trace_event_id_t event_id,
30348     const trace_event_set_t *restrict set,
30349     int *restrict ismember);
30350
```

30351 **DESCRIPTION**

30352 These primitives manipulate sets of trace event types. They operate on data objects addressable  
 30353 by the application, not on the current trace event filter of any trace stream.

30354 The *posix\_trace\_eventset\_add()* and *posix\_trace\_eventset\_del()* functions, respectively, shall add or  
 30355 delete the individual trace event type specified by the value of the argument *event\_id* to or from  
 30356 the trace event type set pointed to by the argument *set*. Adding a trace event type already in the  
 30357 set or deleting a trace event type not in the set shall not be considered an error.

30358 The *posix\_trace\_eventset\_empty()* function shall initialize the trace event type set pointed to by  
 30359 the *set* argument such that all trace event types defined, both system and user, shall be excluded  
 30360 from the set.

30361 The *posix\_trace\_eventset\_fill()* function shall initialize the trace event type set pointed to by the  
 30362 argument *set*, such that the set of trace event types defined by the argument *what* shall be  
 30363 included in the set. The value of the argument *what* shall consist of one of the following values,  
 30364 as defined in the **<trace.h>** header:

30365 **POSIX\_TRACE\_WOPID\_EVENTS**  
 30366 All the process-independent implementation-defined system trace event types are included  
 30367 in the set.

30368 **POSIX\_TRACE\_SYSTEM\_EVENTS** All the implementation-defined system trace event types are  
 30369 included in the set, as are those defined in IEEE Std 1003.1-200x.

30370 **POSIX\_TRACE\_ALL\_EVENTS** All trace event types defined, both system and user, are included  
 30371 in the set.

30372 Applications shall call either *posix\_trace\_eventset\_empty()* or *posix\_trace\_eventset\_fill()* at least  
 30373 once for each object of type **trace\_event\_set\_t** prior to any other use of that object. If such an  
 30374 object is not initialized in this way, but is nonetheless supplied as an argument to any of the  
 30375 *posix\_trace\_eventset\_add()*, *posix\_trace\_eventset\_del()*, or *posix\_trace\_eventset\_ismember()* functions,  
 30376 the results are undefined.

30377 The *posix\_trace\_eventset\_ismember()* function shall test whether the trace event type specified by  
 30378 the value of the argument *event\_id* is a member of the set pointed to by the argument *set*. The  
 30379 value returned in the object pointed to by *ismember* argument is zero if the trace event type  
 30380 identifier is not a member of the set and a value different from zero if it is a member of the set.

**30381 RETURN VALUE**

30382       Upon successful completion, these functions shall return a value of zero. Otherwise, they shall  
30383       return the corresponding error number.

**30384 ERRORS**

30385       These functions may fail if:

30386       [EINVAL]       The value of one of the arguments is invalid.

**30387 EXAMPLES**

30388       None.

**30389 APPLICATION USAGE**

30390       None.

**30391 RATIONALE**

30392       None.

**30393 FUTURE DIRECTIONS**

30394       None.

**30395 SEE ALSO**

30396       *posix\_trace\_set\_filter()*, *posix\_trace\_trid\_eventid\_open()*, the Base Definitions volume of  
30397       IEEE Std 1003.1-200x, <trace.h>

**30398 CHANGE HISTORY**

30399       First released in Issue 6. Derived from IEEE Std 1003.1q-2000.

30400 **NAME**

30401 posix\_trace\_eventtypelist\_getnext\_id, posix\_trace\_eventtypelist\_rewind — iterate over a  
30402 mapping of trace event types (**TRACING**)

30403 **SYNOPSIS**

```
30404 TRC #include <trace.h>
30405 int posix_trace_eventtypelist_getnext_id(trace_id_t trid,
30406     trace_event_id_t *restrict event, int *restrict unavailable);
30407 int posix_trace_eventtypelist_rewind(trace_id_t trid);
30408
```

30409 **DESCRIPTION**

30410 The first time *posix\_trace\_eventtypelist\_getnext\_id()* is called, the function shall return in the  
30411 variable pointed to by *event* the first trace event type identifier of the list of trace events of the  
30412 trace stream identified by the *trid* argument. Successive calls to  
30413 *posix\_trace\_eventtypelist\_getnext\_id()* return in the variable pointed to by *event* the next trace  
30414 event type identifier in that same list. Each time a trace event type identifier is successfully  
30415 written into the variable pointed to by the *event* argument, the variable pointed to by the  
30416 *unavailable* argument shall be set to zero. When no more trace event type identifiers are  
30417 available, and so none is returned, the variable pointed to by the *unavailable* argument shall be  
30418 set to a value different from zero.

30419 The *posix\_trace\_eventtypelist\_rewind()* function shall reset the next trace event type identifier to  
30420 be read to the first trace event type identifier from the list of trace events used in the trace stream  
30421 identified by *trid*.

30422 **RETURN VALUE**

30423 Upon successful completion, these functions shall return a value of zero. Otherwise, they shall  
30424 return the corresponding error number.

30425 The *posix\_trace\_eventtypelist\_getnext\_id()* function stores the trace event type identifier value in  
30426 the object pointed to by *event*, if successful.

30427 **ERRORS**

30428 These functions shall fail if:

30429 [EINVAL] The *trid* argument was not a valid trace stream identifier.

30430 **EXAMPLES**

30431 None.

30432 **APPLICATION USAGE**

30433 None.

30434 **RATIONALE**

30435 None.

30436 **FUTURE DIRECTIONS**

30437 None.

30438 **SEE ALSO**

30439 *posix\_trace\_event()*, *posix\_trace\_getnext\_event()*, *posix\_trace\_trid\_eventid\_open()*, the Base  
30440 Definitions volume of IEEE Std 1003.1-200x, <trace.h>

30441 **CHANGE HISTORY**

30442 First released in Issue 6. Derived from IEEE Std 1003.1q-2000.

30443 IEEE PASC Interpretations 1003.1 #123 and #129 are applied.



30444 **NAME**30445        posix\_trace\_flush — trace stream flush from a process (**TRACING**)30446 **SYNOPSIS**

30447 TRC     #include &lt;sys/types.h&gt;

30448        #include &lt;trace.h&gt;

30449        int posix\_trace\_flush(trace\_id\_t trid);

30450

30451 **DESCRIPTION**30452        Refer to *posix\_trace\_create()*.

30453 **NAME**

30454 posix\_trace\_get\_attr, posix\_trace\_get\_status — retrieve the trace attributes or trace statuses  
 30455 (TRACING)

30456 **SYNOPSIS**

```
30457 TRC #include <trace.h>
30458
30458 int posix_trace_get_attr(trace_id_t trid, trace_attr_t *attr);
30459 int posix_trace_get_status(trace_id_t trid,
30460 struct posix_trace_status_info *statusinfo);
30461
```

30462 **DESCRIPTION**

30463 The *posix\_trace\_get\_attr()* function shall copy the attributes of the active trace stream identified  
 30464 TRL by *trid* into the object pointed to by the *attr* argument. If the Trace Log option is supported, *trid*  
 30465 may represent a pre-recorded trace log.

30466 The *posix\_trace\_get\_status()* function shall return, in the structure pointed to by the *statusinfo*  
 30467 argument, the current trace status for the trace stream identified by the *trid* argument. These  
 30468 status values returned in the structure pointed to by *statusinfo* shall have been appropriately  
 30469 TRL read to ensure that the returned values are consistent. If the Trace Log option is supported and  
 30470 the *trid* argument refers to a pre-recorded trace stream, the status shall be the status of the  
 30471 completed trace stream.

30472 Each time the *posix\_trace\_get\_status()* function is used, the overrun status of the trace stream  
 30473 TRL shall be reset to POSIX\_TRACE\_NO\_OVERRUN immediately after the call completes. If the  
 30474 Trace Log option is supported, the *posix\_trace\_get\_status()* function shall behave the same as  
 30475 when the option is not supported except for the following differences:

- 30476 • If the *trid* argument refers to a trace stream with log, each time the *posix\_trace\_get\_status()*  
 30477 function is used, the log overrun status of the trace stream shall be reset to  
 30478 POSIX\_TRACE\_NO\_OVERRUN and the *flush\_error* status shall be reset to zero immediately  
 30479 after the call completes.

- 30480 • If the *trid* argument refers to a pre-recorded trace stream, the status returned shall be the  
 30481 status of the completed trace stream and the status values of the trace stream shall not be  
 30482 reset.  
 30483

30484 **RETURN VALUE**

30485 Upon successful completion, these functions shall return a value of zero. Otherwise, they shall  
 30486 return the corresponding error number.

30487 The *posix\_trace\_get\_attr()* function stores the trace attributes in the object pointed to by *attr*, if  
 30488 successful.

30489 The *posix\_trace\_get\_status()* function stores the trace status in the object pointed to by *statusinfo*,  
 30490 if successful.

30491 **ERRORS**

30492 These functions shall fail if:

30493 [EINVAL] The trace stream argument *trid* does not correspond to a valid active trace  
 30494 stream or a valid trace log.

30495 **EXAMPLES**

30496 None.

30497 **APPLICATION USAGE**

30498 None.

30499 **RATIONALE**

30500 None.

30501 **FUTURE DIRECTIONS**

30502 None.

30503 **SEE ALSO**30504 *posix\_trace\_attr\_destroy()*, *posix\_trace\_attr\_init()*, *posix\_trace\_create()*, *posix\_trace\_open()*, the Base

30505 Definitions volume of IEEE Std 1003.1-200x, &lt;trace.h&gt;

30506 **CHANGE HISTORY**

30507 First released in Issue 6. Derived from IEEE Std 1003.1q-2000. |

30508 IEEE PASC Interpretation 1003.1 #123 is applied. |

30509 **NAME**

30510 posix\_trace\_get\_filter, posix\_trace\_set\_filter — retrieve and set filter of an initialized trace  
30511 stream (**TRACING**)

30512 **SYNOPSIS**

```
30513 TRC TEF #include <trace.h>
30514
30514 int posix_trace_get_filter(trace_id_t trid, trace_event_set_t *set);
30515 int posix_trace_set_filter(trace_id_t trid,
30516     const trace_event_set_t *set, int how);
30517
```

30518 **DESCRIPTION**

30519 The *posix\_trace\_get\_filter()* function shall retrieve, into the argument pointed to by *set*, the actual  
30520 trace event filter from the trace stream specified by *trid*.

30521 The *posix\_trace\_set\_filter()* function shall change the set of filtered trace event types after a trace  
30522 stream identified by the *trid* argument is created. This function may be called prior to starting  
30523 the trace stream, or while the trace stream is active. By default, if no call is made to  
30524 *posix\_trace\_set\_filter()*, all trace events shall be recorded (that is, none of the trace event types are  
30525 filtered out).

30526 If this function is called while the trace is in progress, a special system trace event,  
30527 POSIX\_TRACE\_FILTER, shall be recorded in the trace indicating both the old and the new sets  
30528 of filtered trace event types (see Table 2-4 (on page 528) and Table 2-6 (on page 529)).

30529 If the *posix\_trace\_set\_filter()* function is interrupted by a signal, an error shall be returned and the  
30530 filter shall not be changed. In this case, the state of the trace stream shall not be changed.

30531 The value of the argument *how* indicates the manner in which the set is to be changed and shall  
30532 have one of the following values, as defined in the **<trace.h>** header:

30533 **POSIX\_TRACE\_SET\_EVENTSET**  
30534 The resulting set of trace event types to be filtered shall be the trace event type set pointed  
30535 to by the argument *set*.

30536 **POSIX\_TRACE\_ADD\_EVENTSET**  
30537 The resulting set of trace event types to be filtered shall be the union of the current set and  
30538 the trace event type set pointed to by the argument *set*.

30539 **POSIX\_TRACE\_SUB\_EVENTSET**  
30540 The resulting set of trace event types to be filtered shall be all trace event types in the  
30541 current set that are not in the set pointed to by the argument *set*; that is, remove each  
30542 element of the specified set from the current filter.

30543 **RETURN VALUE**

30544 Upon successful completion, these functions shall return a value of zero. Otherwise, they shall  
30545 return the corresponding error number.

30546 The *posix\_trace\_get\_filter()* function stores the set of filtered trace event types in *set*, if successful.

30547 **ERRORS**

30548 These functions shall fail if:

30549 [EINVAL] The value of the *trid* argument does not correspond to an active trace stream  
30550 or the value of the argument pointed to by *set* is invalid.

30551 [EINTR] The operation was interrupted by a signal.

30552 **EXAMPLES**

30553 None.

30554 **APPLICATION USAGE**

30555 None.

30556 **RATIONALE**

30557 None.

30558 **FUTURE DIRECTIONS**

30559 None.

30560 **SEE ALSO**30561 *posix\_trace\_eventset\_add()*, the Base Definitions volume of IEEE Std 1003.1-200x, <trace.h>30562 **CHANGE HISTORY**

30563 First released in Issue 6. Derived from IEEE Std 1003.1q-2000. |

30564 IEEE PASC Interpretation 1003.1 #123 is applied. |

30565 **NAME**

30566        posix\_trace\_get\_status — retrieve the trace statuses (**TRACING**)

30567 **SYNOPSIS**

30568 TRC     #include <trace.h>

```
30569        int posix_trace_get_status(trace_id_t trid,  
30570                                  struct posix_trace_status_info *statusinfo);
```

30571

30572 **DESCRIPTION**

30573        Refer to *posix\_trace\_get\_attr()*.

## 30574 NAME

30575 posix\_trace\_getnext\_event, posix\_trace\_timedgetnext\_event, posix\_trace\_trygetnext\_event —  
 30576 retrieve a trace event (**TRACING**)

## 30577 SYNOPSIS

```
30578 TRC #include <sys/types.h>
30579 #include <trace.h>

30580 int posix_trace_getnext_event(trace_id_t trid,
30581 struct posix_trace_event_info *restrict event,
30582 void *restrict data, size_t num_bytes,
30583 size_t *restrict data_len, int *restrict unavailable);
30584 TRC TMO int posix_trace_timedgetnext_event(trace_id_t trid,
30585 struct posix_trace_event_info *restrict event,
30586 void *restrict data, size_t num_bytes,
30587 size_t *restrict data_len, int *restrict unavailable,
30588 const struct timespec *restrict abs_timeout);
30589 TRC int posix_trace_trygetnext_event(trace_id_t trid,
30590 struct posix_trace_event_info *restrict event,
30591 void *restrict data, size_t num_bytes,
30592 size_t *restrict data_len, int *restrict unavailable);
30593
```

## 30594 DESCRIPTION

30595 The *posix\_trace\_getnext\_event()* function shall report a recorded trace event either from an active |  
 30596 TRL trace stream without log or a pre-recorded trace stream identified by the *trid* argument. The |  
 30597 *posix\_trace\_trygetnext\_event()* function shall report a recorded trace event from an active trace |  
 30598 stream without log identified by the *trid* argument.

30599 The trace event information associated with the recorded trace event shall be copied by the  
 30600 function into the structure pointed to by the argument *event* and the data associated with the  
 30601 trace event shall be copied into the buffer pointed to by the *data* argument.

30602 The *posix\_trace\_getnext\_event()* function shall block if the *trid* argument identifies an active trace  
 30603 stream and there is currently no trace event ready to be retrieved. When returning, if a recorded  
 30604 trace event was reported, the variable pointed to by the *unavailable* argument shall be set to zero.  
 30605 Otherwise, the variable pointed to by the *unavailable* argument shall be set to a value different  
 30606 from zero.

30607 TMO The *posix\_trace\_timedgetnext\_event()* function shall attempt to get another trace event from an |  
 30608 active trace stream without log, as in the *posix\_trace\_getnext\_event()* function. However, if no |  
 30609 trace event is available from the trace stream, the implied wait shall be terminated when the |  
 30610 timeout specified by the argument *abs\_timeout* expires, and the function shall return the error |  
 30611 [ETIMEDOUT].

30612 The timeout shall expire when the absolute time specified by *abs\_timeout* passes, as measured by |  
 30613 the clock upon which timeouts are based (that is, when the value of that clock equals or exceeds  
 30614 *abs\_timeout*), or if the absolute time specified by *abs\_timeout* has already passed at the time of the  
 30615 call.

30616 TMO TMR If the Timers option is supported, the timeout shall be based on the CLOCK\_REALTIME clock; |  
 30617 if the Timers option is not supported, the timeout shall be based on the system clock as returned |  
 30618 by the *time()* function. The resolution of the timeout shall be the resolution of the clock on which |  
 30619 it is based. The **timespec** data type is defined in the **<time.h>** header. |

30620 TMO Under no circumstance shall the function fail with a timeout if a trace event is immediately |  
 30621 available from the trace stream. The validity of the *abs\_timeout* argument need not be checked if

30622 a trace event is immediately available from the trace stream.

30623 The behavior of this function for a pre-recorded trace stream is unspecified.

30624 TRL The *posix\_trace\_trygetnext\_event()* function shall not block. This function shall return an error if  
 30625 the *trid* argument identifies a pre-recorded trace stream. If a recorded trace event was reported,  
 30626 the variable pointed to by the *unavailable* argument shall be set to zero. Otherwise, if no trace  
 30627 event was reported, the variable pointed to by the *unavailable* argument shall be set to a value  
 30628 different from zero.

30629 The argument *num\_bytes* shall be the size of the buffer pointed to by the *data* argument. The  
 30630 argument *data\_len* reports to the application the length in bytes of the data record just  
 30631 transferred. If *num\_bytes* is greater than or equal to the size of the data associated with the trace  
 30632 event pointed to by the *event* argument, all the recorded data shall be transferred. In this case, the  
 30633 *truncation-status* member of the trace event structure shall be either  
 30634 POSIX\_TRACE\_NOT\_TRUNCATED, if the trace event data was recorded without truncation  
 30635 while tracing, or POSIX\_TRACE\_TRUNCATED\_RECORD, if the trace event data was truncated  
 30636 when it was recorded. If the *num\_bytes* argument is less than the length of recorded trace event  
 30637 data, the data transferred shall be truncated to a length of *num\_bytes*, the value stored in the  
 30638 variable pointed to by *data\_len* shall be equal to *num\_bytes*, and the *truncation-status* member of  
 30639 the *event* structure argument shall be set to POSIX\_TRACE\_TRUNCATED\_READ (see the  
 30640 *posix\_trace\_event\_info()* function).

30641 The report of a trace event shall be sequential starting from the oldest recorded trace event. Trace  
 30642 events shall be reported in the order in which they were generated, up to an implementation-  
 30643 defined time resolution that causes the ordering of trace events occurring very close to each  
 30644 other to be unknown. Once reported, a trace event cannot be reported again from an active trace  
 30645 stream. Once a trace event is reported from an active trace stream without log, the trace stream  
 30646 shall make the resources associated with that trace event available to record future generated  
 30647 trace events.

30648 **RETURN VALUE**

30649 Upon successful completion, these functions shall return a value of zero. Otherwise, they shall  
 30650 return the corresponding error number.

30651 If successful, these functions store:

- 30652 • The recorded trace event in the object pointed to by *event*
- 30653 • The trace event information associated with the recorded trace event in the object pointed to  
 30654 by *data*
- 30655 • The length of this trace event information in the object pointed to by *data\_len*
- 30656 • The value of zero in the object pointed to by *unavailable*

30657 **ERRORS**

30658 These functions shall fail if:

30659 [EINVAL] The trace stream identifier argument *trid* is invalid.

30660 The *posix\_trace\_getnext\_event()* and *posix\_trace\_timedgetnext\_event()* functions shall fail if:

30661 [EINTR] The operation was interrupted by a signal, and so the call had no effect.

30662 The *posix\_trace\_trygetnext\_event()* function shall fail if:

30663 [EINVAL] The trace stream identifier argument *trid* does not correspond to an active  
 30664 trace stream.



30665 TMO The *posix\_trace\_timedgetnext\_event()* function shall fail if: |

30666 [EINVAL] There is no trace event immediately available from the trace stream, and the  
30667 *timeout* argument is invalid.

30668 [ETIMEDOUT] No trace event was available from the trace stream before the specified  
30669 *timeout* *timeout* expired.  
30670

30671 **EXAMPLES**

30672 None.

30673 **APPLICATION USAGE**

30674 None.

30675 **RATIONALE**

30676 None.

30677 **FUTURE DIRECTIONS**

30678 None.

30679 **SEE ALSO**

30680 *posix\_trace\_create()*, **posix\_trace\_event\_info Structure**, *posix\_trace\_open()*, the Base Definitions  
30681 volume of IEEE Std 1003.1-200x, <sys/types.h>, <trace.h>

30682 **CHANGE HISTORY**

30683 First released in Issue 6. Derived from IEEE Std 1003.1q-2000. |

30684 IEEE PASC Interpretation 1003.1 #123 is applied. |

30685 **NAME**

30686        posix\_trace\_open — trace log management (**TRACING**)

30687 **SYNOPSIS**

30688 TCT TRL   #include <trace.h>

30689        int posix\_trace\_open(int *file\_desc*, trace\_id\_t \*trid);

30690

30691 **DESCRIPTION**

30692        Refer to *posix\_trace\_close()*.

30693 **NAME**30694        posix\_trace\_rewind — trace log management (**TRACING**)30695 **SYNOPSIS**30696 TCT TRL `#include <trace.h>`30697        `int posix_trace_rewind(trace_id_t trid);`

30698

30699 **DESCRIPTION**30700        Refer to *posix\_trace\_close()*.

30701 **NAME**

30702        posix\_trace\_set\_filter — set filter of an initialized trace stream (**TRACING**)

30703 **SYNOPSIS**

30704 TRC TEF #include <trace.h>

```
30705     int posix_trace_set_filter(trace_id_t trid,  
30706                             const trace_event_set_t *set, int how);
```

30707

30708 **DESCRIPTION**

30709        Refer to *posix\_trace\_get\_filter()*.

30710 **NAME**

30711        posix\_trace\_shutdown — trace stream shutdown from a process (**TRACING**)

30712 **SYNOPSIS**

30713 TRC     #include <sys/types.h>

30714        #include <trace.h>

30715        int posix\_trace\_shutdown(trace\_id\_t *trid*);

30716

30717 **DESCRIPTION**

30718        Refer to *posix\_trace\_create()*.

30719 **NAME**

30720 posix\_trace\_start, posix\_trace\_stop — trace start and stop (**TRACING**)

30721 **SYNOPSIS**

30722 TRC `#include <trace.h>`

30723 `int posix_trace_start(trace_id_t trid);`

30724 `int posix_trace_stop (trace_id_t trid);`

30725

30726 **DESCRIPTION**

30727 The *posix\_trace\_start()* and *posix\_trace\_stop()* functions, respectively, shall start and stop the  
30728 trace stream identified by the argument *trid*.

30729 The effect of calling the *posix\_trace\_start()* function shall be recorded in the trace stream as the  
30730 POSIX\_TRACE\_START system trace event and the status of the trace stream shall become  
30731 POSIX\_TRACE\_RUNNING. If the trace stream is in progress when this function is called, the  
30732 POSIX\_TRACE\_START system trace event shall not be recorded and the trace stream shall  
30733 continue to run. If the trace stream is full, the POSIX\_TRACE\_START system trace event shall  
30734 not be recorded and the status of the trace stream shall not be changed.

30735 The effect of calling the *posix\_trace\_stop()* function shall be recorded in the trace stream as the  
30736 POSIX\_TRACE\_STOP system trace event and the status of the trace stream shall become  
30737 POSIX\_TRACE\_SUSPENDED. If the trace stream is suspended when this function is called, the  
30738 POSIX\_TRACE\_STOP system trace event shall not be recorded and the trace stream shall remain  
30739 suspended. If the trace stream is full, the POSIX\_TRACE\_STOP system trace event shall not be  
30740 recorded and the status of the trace stream shall not be changed.

30741 **RETURN VALUE**

30742 Upon successful completion, these functions shall return a value of zero. Otherwise, they shall  
30743 return the corresponding error number.

30744 **ERRORS**

30745 These functions shall fail if:

30746 [EINVAL] The value of the argument *trid* does not correspond to an active trace stream  
30747 and thus no trace stream was started or stopped.

30748 [EINTR] The operation was interrupted by a signal and thus the trace stream was not  
30749 necessarily started or stopped.

30750 **EXAMPLES**

30751 None.

30752 **APPLICATION USAGE**

30753 None.

30754 **RATIONALE**

30755 None.

30756 **FUTURE DIRECTIONS**

30757 None.

30758 **SEE ALSO**

30759 *posix\_trace\_create()*, the Base Definitions volume of IEEE Std 1003.1-200x, `<trace.h>`

30760 **CHANGE HISTORY**

30761 First released in Issue 6. Derived from IEEE Std 1003.1q-2000.

30762 IEEE PASC Interpretation 1003.1 #123 is applied.

30763 **NAME**

30764        posix\_trace\_timedgetnext\_event, — retrieve a trace event (**TRACING**)

30765 **SYNOPSIS**

```
30766 TRC TMO #include <sys/types.h>
30767           #include <trace.h>
```

```
30768        int posix_trace_timedgetnext_event(trace_id_t trid,
30769                                          struct posix_trace_event_info *restrict event,
30770                                          void *restrict data, size_t num_bytes,
30771                                          size_t *restrict data_len, int *restrict unavailable,
30772                                          const struct timespec *restrict abs_timeout);
30773
```

30774 **DESCRIPTION**

30775        Refer to *posix\_trace\_getnext\_event()*.



30776 **NAME**30777        posix\_trace\_trid\_eventid\_open — manipulate trace event type identifier (**TRACING**)30778 **SYNOPSIS**

30779 TRC TEF #include &lt;trace.h&gt;

```
30780 int posix_trace_trid_eventid_open(trace_id_t trid,  
30781     const char *restrict event_name,  
30782     trace_event_id_t *restrict event);
```

30783

30784 **DESCRIPTION**30785        Refer to *posix\_trace\_eventid\_equal()*.

30786 **NAME**

30787        posix\_trace\_trygetnext\_event — retrieve a trace event (**TRACING**)

30788 **SYNOPSIS**

30789 TRC     #include <sys/types.h>

30790        #include <trace.h>

```
30791        int posix_trace_trygetnext_event(trace_id_t trid,  
30792                                        struct posix_trace_event_info *restrict event,  
30793                                        void *restrict data, size_t num_bytes,  
30794                                        size_t *restrict data_len, int *restrict unavailable);  
30795
```

30796 **DESCRIPTION**

30797        Refer to *posix\_trace\_getnext\_event()*.

30798 **NAME**30799 `posix_typed_mem_get_info` — query typed memory information (**ADVANCED REALTIME**)30800 **SYNOPSIS**30801 `TYM` `#include <sys/mman.h>`30802 `int posix_typed_mem_get_info(int fildes,`  
30803 `struct posix_typed_mem_info *info);`

30804

30805 **DESCRIPTION**

30806 The `posix_typed_mem_get_info()` function shall return, in the `posix_tmi_length` field of the  
 30807 **posix\_typed\_mem\_info** structure pointed to by `info`, the maximum length which may be  
 30808 successfully allocated by the typed memory object designated by `fildes`. This maximum length  
 30809 shall take into account the flag `POSIX_TYPED_MEM_ALLOCATE` or  
 30810 `POSIX_TYPED_MEM_ALLOCATE_CONTIG` specified when the typed memory object  
 30811 represented by `fildes` was opened. The maximum length is dynamic; therefore, the value returned  
 30812 is valid only while the current mapping of the corresponding typed memory pool remains  
 30813 unchanged.

30814 If `fildes` represents a typed memory object opened with neither the  
 30815 `POSIX_TYPED_MEM_ALLOCATE` flag nor the `POSIX_TYPED_MEM_ALLOCATE_CONTIG`  
 30816 flag specified, the returned value of `info->posix_tmi_length` is unspecified.

30817 The `posix_typed_mem_get_info()` function may return additional implementation-defined  
 30818 information in other fields of the **posix\_typed\_mem\_info** structure pointed to by `info`.

30819 If the memory object specified by `fildes` is not a typed memory object, then the behavior of this  
 30820 function is undefined.

30821 **RETURN VALUE**

30822 Upon successful completion, the `posix_typed_mem_get_info()` function shall return zero;  
 30823 otherwise, the corresponding error status value shall be returned.

30824 **ERRORS**

30825 The `posix_typed_mem_get_info()` function shall fail if:

30826 [EBADF] The `fildes` argument is not a valid open file descriptor.

30827 [ENODEV] The `fildes` argument is not connected to a memory object supported by this  
 30828 function.

30829 This function shall not return an error code of [EINTR].

30830 **EXAMPLES**

30831 None.

30832 **APPLICATION USAGE**

30833 None.

30834 **RATIONALE**

30835 An application that needs to allocate a block of typed memory with length dependent upon the  
 30836 amount of memory currently available must either query the typed memory object to obtain the  
 30837 amount available, or repeatedly invoke `mmap()` attempting to guess an appropriate length.  
 30838 While the latter method is existing practice with `malloc()`, it is awkward and imprecise. The  
 30839 `posix_typed_mem_get_info()` function allows an application to immediately determine available  
 30840 memory. This is particularly important for typed memory objects that may in some cases be  
 30841 scarce resources. Note that when a typed memory pool is a shared resource, some form of  
 30842 mutual exclusion or synchronization may be required while typed memory is being queried and

30843 allocated to prevent race conditions.

30844 The existing *fstat()* function is not suitable for this purpose. We realize that implementations  
30845 may wish to provide other attributes of typed memory objects (for example, alignment  
30846 requirements, page size, and so on). The *fstat()* function returns a structure which is not  
30847 extensible and, furthermore, contains substantial information that is inappropriate for typed  
30848 memory objects.

30849 **FUTURE DIRECTIONS**

30850 None.

30851 **SEE ALSO**

30852 *fstat()*, *mmap()*, *posix\_typed\_mem\_open()*, the Base Definitions volume of IEEE Std 1003.1-200x,  
30853 `<sys/mman.h>`

30854 **CHANGE HISTORY**

30855 First released in Issue 6. Derived from IEEE Std 1003.1j-2000.

30856 **NAME**30857 posix\_typed\_mem\_open — open a typed memory object (**ADVANCED REALTIME**)30858 **SYNOPSIS**

30859 TYM #include &lt;sys/mman.h&gt;

30860 int posix\_typed\_mem\_open(const char \*name, int oflag, int tflag);

30861

30862 **DESCRIPTION**

30863 The *posix\_typed\_mem\_open()* function shall establish a connection between the typed memory  
 30864 object specified by the string pointed to by *name* and a file descriptor. It shall create an open file  
 30865 description that refers to the typed memory object and a file descriptor that refers to that open  
 30866 file description. The file descriptor is used by other functions to refer to that typed memory  
 30867 object. It is unspecified whether the name appears in the file system and is visible to other  
 30868 functions that take pathnames as arguments. The *name* argument shall conform to the  
 30869 construction rules for a pathname. If *name* begins with the slash character, then processes calling  
 30870 *posix\_typed\_mem\_open()* with the same value of *name* shall refer to the same typed memory  
 30871 object. If *name* does not begin with the slash character, the effect is implementation-defined. The  
 30872 interpretation of slash characters other than the leading slash character in *name* is  
 30873 implementation-defined.

30874 Each typed memory object supported in a system shall be identified by a name which specifies  
 30875 not only its associated typed memory pool, but also the path or port by which it is accessed. That  
 30876 is, the same typed memory pool accessed via several different ports shall have several different  
 30877 corresponding names. The binding between names and typed memory objects is established in  
 30878 an implementation-defined manner. Unlike shared memory objects, there is no way within  
 30879 IEEE Std 1003.1-200x for a program to create a typed memory object.

30880 The value of *tflag* shall determine how the typed memory object behaves when subsequently  
 30881 mapped by calls to *mmap()*. At most, one of the following flags defined in <sys/mman.h> may  
 30882 be specified:

30883 POSIX\_TYPED\_MEM\_ALLOCATE

30884 Allocate on *mmap()*.

30885 POSIX\_TYPED\_MEM\_ALLOCATE\_CONTIG

30886 Allocate contiguously on *mmap()*.

30887 POSIX\_TYPED\_MEM\_MAP\_ALLOCATABLE

30888 Map on *mmap()*, without affecting allocatability.

30889 If *tflag* has the flag POSIX\_TYPED\_MEM\_ALLOCATE specified, any subsequent call to *mmap()*  
 30890 using the returned file descriptor shall result in allocation and mapping of typed memory from  
 30891 the specified typed memory pool. The allocated memory may be a contiguous previously  
 30892 unallocated area of the typed memory pool or several non-contiguous previously unallocated  
 30893 areas (mapped to a contiguous portion of the process address space). If *tflag* has the flag  
 30894 POSIX\_TYPED\_MEM\_ALLOCATE\_CONTIG specified, any subsequent call to *mmap()* using the  
 30895 returned file descriptor shall result in allocation and mapping of a single contiguous previously  
 30896 unallocated area of the typed memory pool (also mapped to a contiguous portion of the process  
 30897 address space). If *tflag* has none of the flags POSIX\_TYPED\_MEM\_ALLOCATE or  
 30898 POSIX\_TYPED\_MEM\_ALLOCATE\_CONTIG specified, any subsequent call to *mmap()* using the  
 30899 returned file descriptor shall map an application-chosen area from the specified typed memory  
 30900 pool such that this mapped area becomes unavailable for allocation until unmapped by all  
 30901 processes. If *tflag* has the flag POSIX\_TYPED\_MEM\_MAP\_ALLOCATABLE specified, any  
 30902 subsequent call to *mmap()* using the returned file descriptor shall map an application-chosen  
 30903 area from the specified typed memory pool without an effect on the availability of that area for

30904 allocation; that is, mapping such an object leaves each byte of the mapped area unallocated if it  
 30905 was unallocated prior to the mapping or allocated if it was allocated prior to the mapping. The  
 30906 appropriate privilege to specify the POSIX\_TYPED\_MEM\_MAP\_ALLOCATABLE flag is  
 30907 implementation-defined.

30908 If successful, *posix\_typed\_mem\_open()* shall return a file descriptor for the typed memory object  
 30909 that is the lowest numbered file descriptor not currently open for that process. The open file  
 30910 description is new, and therefore the file descriptor shall not share it with any other processes. It  
 30911 is unspecified whether the file offset is set. The FD\_CLOEXEC file descriptor flag associated  
 30912 with the new file descriptor shall be cleared.

30913 The behavior of *msync()*, *ftruncate()*, and all file operations other than *mmap()*,  
 30914 *posix\_mem\_offset()*, *posix\_typed\_mem\_get\_info()*, *fstat()*, *dup()*, *dup2()*, and *close()*, is unspecified  
 30915 when passed a file descriptor connected to a typed memory object by this function.

30916 The file status flags of the open file description shall be set according to the value of *oflag*.  
 30917 Applications shall specify exactly one of the three access mode values described below and  
 30918 defined in the `<fcntl.h>` header, as the value of *oflag*.

- 30919 O\_RDONLY Open for read access only.
- 30920 O\_WRONLY Open for write access only.
- 30921 O\_RDWR Open for read or write access.

30922 **RETURN VALUE**

30923 Upon successful completion, the *posix\_typed\_mem\_open()* function shall return a non-negative  
 30924 integer representing the lowest numbered unused file descriptor. Otherwise, it shall return -1  
 30925 and set *errno* to indicate the error.

30926 **ERRORS**

30927 The *posix\_typed\_mem\_open()* function shall fail if:

- 30928 [EACCES] The typed memory object exists and the permissions specified by *oflag* are  
 30929 denied.
- 30930 [EINTR] The *posix\_typed\_mem\_open()* operation was interrupted by a signal.
- 30931 [EINVAL] The flags specified in *tflag* are invalid (more than one of  
 30932 POSIX\_TYPED\_MEM\_ALLOCATE,  
 30933 POSIX\_TYPED\_MEM\_ALLOCATE\_CONTIG, or  
 30934 POSIX\_TYPED\_MEM\_MAP\_ALLOCATABLE is specified).
- 30935 [EMFILE] Too many file descriptors are currently in use by this process.
- 30936 [ENAMETOOLONG]  
 30937 The length of the *name* argument exceeds {PATH\_MAX} or a pathname  
 30938 component is longer than {NAME\_MAX}.
- 30939 [ENFILE] Too many file descriptors are currently open in the system.
- 30940 [ENOENT] The named typed memory object does not exist.
- 30941 [EPERM] The caller lacks the appropriate privilege to specify the flag  
 30942 POSIX\_TYPED\_MEM\_MAP\_ALLOCATABLE in argument *tflag*.

30943 **EXAMPLES**

30944 None.

30945 **APPLICATION USAGE**

30946 None.

30947 **RATIONALE**

30948 None.

30949 **FUTURE DIRECTIONS**

30950 None.

30951 **SEE ALSO**

30952 *close()*, *dup()*, *exec*, *fcntl()*, *fstat()*, *ftruncate()*, *mmap()*, *msync()*, *posix\_mem\_offset()*,  
30953 *posix\_typed\_mem\_get\_info()*, *umask()*, the Base Definitions volume of IEEE Std 1003.1-200x,  
30954 *<fcntl.h>*, *<sys/mman.h>*

30955 **CHANGE HISTORY**

30956 First released in Issue 6. Derived from IEEE Std 1003.1j-2000.

30957 **NAME**

30958 pow, powf, powl — power function

30959 **SYNOPSIS**

30960 #include &lt;math.h&gt;

30961 double pow(double x, double y);

30962 float powf(float x, float y);

30963 long double powl(long double x, long double y);

30964 **DESCRIPTION**

30965 **CX** The functionality described on this reference page is aligned with the ISO C standard. Any  
 30966 conflict between the requirements described here and the ISO C standard is unintentional. This  
 30967 volume of IEEE Std 1003.1-200x defers to the ISO C standard.

30968 These functions shall compute the value of  $x$  raised to the power  $y$ ,  $x^y$ . If  $x$  is negative, the  
 30969 application shall ensure that  $y$  is an integer value.

30970 An application wishing to check for error situations should set *errno* to zero and call  
 30971 *feclearexcept*(FE\_ALL\_EXCEPT) before calling these functions. On return, if *errno* is non-zero or  
 30972 *fetestexcept*(FE\_INVALID | FE\_DIVBYZERO | FE\_OVERFLOW | FE\_UNDERFLOW) is non-  
 30973 zero, an error has occurred.

30974 **RETURN VALUE**30975 Upon successful completion, these functions shall return the value of  $x$  raised to the power  $y$ .

30976 **MX** For finite values of  $x < 0$ , and finite non-integer values of  $y$ , a domain error shall occur and either  
 30977 a NaN (if representable), or an implementation-defined value shall be returned.

30978 If the correct value would cause overflow, a range error shall occur and *pow()*, *powf()*, and  
 30979 *powl()* shall return HUGE\_VAL, HUGE\_VALF, and HUGE\_VALL, respectively.

30980 If the correct value would cause underflow, and is not representable, a range error may occur,  
 30981 **MX** and either 0.0 (if supported), or an implementation-defined value shall be returned.

30982 **MX** If  $x$  or  $y$  is a NaN, a NaN shall be returned (unless specified elsewhere in this description).30983 For any value of  $y$  (including NaN), if  $x$  is +1, 1.0 shall be returned.30984 For any value of  $x$  (including NaN), if  $y$  is  $\pm 0$ , 1.0 shall be returned.30985 For any odd integer value of  $y > 0$ , if  $x$  is  $\pm 0$ ,  $\pm 0$  shall be returned.30986 For  $y > 0$  and not an odd integer, if  $x$  is  $\pm 0$ , +0 shall be returned.30987 If  $x$  is  $-1$ , and  $y$  is  $\pm \text{Inf}$ , 1.0 shall be returned.30988 For  $|x| < 1$ , if  $y$  is  $-\text{Inf}$ ,  $+\text{Inf}$  shall be returned.30989 For  $|x| > 1$ , if  $y$  is  $-\text{Inf}$ , +0 shall be returned.30990 For  $|x| < 1$ , if  $y$  is  $+\text{Inf}$ , +0 shall be returned.30991 For  $|x| > 1$ , if  $y$  is  $+\text{Inf}$ ,  $+\text{Inf}$  shall be returned.30992 For  $y$  an odd integer  $< 0$ , if  $x$  is  $-\text{Inf}$ ,  $-0$  shall be returned.30993 For  $y < 0$  and not an odd integer, if  $x$  is  $-\text{Inf}$ , +0 shall be returned.30994 For  $y$  an odd integer  $> 0$ , if  $x$  is  $-\text{Inf}$ ,  $-\text{Inf}$  shall be returned.30995 For  $y > 0$  and not an odd integer, if  $x$  is  $-\text{Inf}$ ,  $+\text{Inf}$  shall be returned.



30996 For  $y < 0$ , if  $x$  is  $+\text{Inf}$ ,  $+0$  shall be returned.

30997 For  $y > 0$ , if  $x$  is  $+\text{Inf}$ ,  $+\text{Inf}$  shall be returned.

30998 For  $y$  an odd integer  $< 0$ , if  $x$  is  $\pm 0$ , a pole error shall occur and  $\pm\text{HUGE\_VAL}$ ,  $\pm\text{HUGE\_VALF}$ , and  
 30999  $\pm\text{HUGE\_VALL}$  shall be returned for  $\text{pow}()$ ,  $\text{powf}()$ , and  $\text{powl}()$ , respectively.

31000 For  $y < 0$  and not an odd integer, if  $x$  is  $\pm 0$ , a pole error shall occur and  $\text{HUGE\_VAL}$ ,  
 31001  $\text{HUGE\_VALF}$ , and  $\text{HUGE\_VALL}$  shall be returned for  $\text{pow}()$ ,  $\text{powf}()$ , and  $\text{powl}()$ , respectively.

31002 If the correct value would cause underflow, and is representable, a range error may occur and  
 31003 the correct value shall be returned.

**31004 ERRORS**

31005 These functions shall fail if:

31006 Domain Error The value of  $x$  is negative and  $y$  is a finite non-integer.

31007 If the integer expression ( $\text{math\_errhandling}$  &  $\text{MATH\_ERRNO}$ ) is non-zero, |  
 31008 then *errno* shall be set to [EDOM]. If the integer expression ( $\text{math\_errhandling}$  |  
 31009 &  $\text{MATH\_ERREXCEPT}$ ) is non-zero, then the invalid floating-point exception |  
 31010 shall be raised. |

31011 MX Pole Error The value of  $x$  is zero and  $y$  is negative.

31012 If the integer expression ( $\text{math\_errhandling}$  &  $\text{MATH\_ERRNO}$ ) is non-zero, |  
 31013 then *errno* shall be set to [ERANGE]. If the integer expression |  
 31014 ( $\text{math\_errhandling}$  &  $\text{MATH\_ERREXCEPT}$ ) is non-zero, then the divide-by- |  
 31015 zero floating-point exception shall be raised. |

31016 Range Error The result overflows.

31017 If the integer expression ( $\text{math\_errhandling}$  &  $\text{MATH\_ERRNO}$ ) is non-zero, |  
 31018 then *errno* shall be set to [ERANGE]. If the integer expression |  
 31019 ( $\text{math\_errhandling}$  &  $\text{MATH\_ERREXCEPT}$ ) is non-zero, then the overflow |  
 31020 floating-point exception shall be raised. |

31021 These functions may fail if:

31022 Range Error The result underflows.

31023 If the integer expression ( $\text{math\_errhandling}$  &  $\text{MATH\_ERRNO}$ ) is non-zero, |  
 31024 then *errno* shall be set to [ERANGE]. If the integer expression |  
 31025 ( $\text{math\_errhandling}$  &  $\text{MATH\_ERREXCEPT}$ ) is non-zero, then the underflow |  
 31026 floating-point exception shall be raised. |

**31027 EXAMPLES**

31028 None.

**31029 APPLICATION USAGE**

31030 On error, the expressions ( $\text{math\_errhandling}$  &  $\text{MATH\_ERRNO}$ ) and ( $\text{math\_errhandling}$  &  
 31031  $\text{MATH\_ERREXCEPT}$ ) are independent of each other, but at least one of them must be non-zero.

**31032 RATIONALE**

31033 None.

**31034 FUTURE DIRECTIONS**

31035 None.

31036 **SEE ALSO**

31037 *exp()*, *feclearexcept()*, *fetestexcept()*, *isnan()*, the Base Definitions volume of IEEE Std 1003.1-200x, |  
31038 Section 4.18, Treatment of Error Conditions for Mathematical Functions, <math.h> |

31039 **CHANGE HISTORY**

31040 First released in Issue 1. Derived from Issue 1 of the SVID.

31041 **Issue 5**

31042 The DESCRIPTION is updated to indicate how an application should check for an error. This  
31043 text was previously published in the APPLICATION USAGE section.

31044 **Issue 6**

31045 The DESCRIPTION is updated to avoid use of the term “must” for application requirements.

31046 The *powf()* and *powl()* functions are added for alignment with the ISO/IEC 9899:1999 standard.

31047 The DESCRIPTION, RETURN VALUE, ERRORS, and APPLICATION USAGE sections are  
31048 revised to align with the ISO/IEC 9899:1999 standard.

31049 IEC 60559:1989 standard floating-point extensions over the ISO/IEC 9899:1999 standard are  
31050 marked.

31051 **NAME**

31052        pread — read from a file

31053 **SYNOPSIS**

31054 xSI        #include &lt;unistd.h&gt;

31055        ssize\_t pread(int *fd*, void \**buf*, size\_t *nbyte*, off\_t *offset*);

31056

31057 **DESCRIPTION**31058        Refer to *read()*.

31059 **NAME**

31060       printf — print formatted output

31061 **SYNOPSIS**

31062       #include <stdio.h>

31063       int printf(const char \*restrict *format*, ...);

31064 **DESCRIPTION**

31065       Refer to *fprintf()*.

31066 **NAME**

31067 pselect, select — synchronous I/O multiplexing

31068 **SYNOPSIS**

```

31069 #include <sys/select.h>

31070 int pselect(int nfds, fd_set *restrict readfds,
31071            fd_set *restrict writefds, fd_set *restrict errorfds,
31072            const struct timespec *restrict timeout,
31073            const sigset_t *restrict sigmask);
31074 int select(int nfds, fd_set *restrict readfds,
31075           fd_set *restrict writefds, fd_set *restrict errorfds,
31076           struct timeval *restrict timeout);
31077 void FD_CLR(int fd, fd_set *fdset);
31078 int FD_ISSET(int fd, fd_set *fdset);
31079 void FD_SET(int fd, fd_set *fdset);
31080 void FD_ZERO(fd_set *fdset);

```

31081 **DESCRIPTION**

31082 The *pselect()* function shall examine the file descriptor sets whose addresses are passed in the  
 31083 *readfds*, *writefds*, and *errorfds* parameters to see if some of their descriptors are ready for reading,  
 31084 are ready for writing, or have an exceptional condition pending, respectively.

31085 The *select()* function shall be equivalent to the *pselect()* function, except as follows:

- 31086 • For the *select()* function, the timeout period is given in seconds and microseconds in an  
 31087 argument of type **struct timeval**, whereas for the *pselect()* function the timeout period is  
 31088 given in seconds and nanoseconds in an argument of type **struct timespec**.
- 31089 • The *select()* function has no *sigmask* argument; it shall behave as *pselect()* does when *sigmask*  
 31090 is a null pointer.
- 31091 • Upon successful completion, the *select()* function may modify the object pointed to by the  
 31092 *timeout* argument.

31093 The *pselect()* and *select()* functions shall support regular files, terminal and pseudo-terminal  
 31094 XSR devices, **STREAMS**-based files, FIFOs, pipes, and sockets. The behavior of *pselect()* and *select()*  
 31095 on file descriptors that refer to other types of file is unspecified.

31096 The *nfds* argument specifies the range of descriptors to be tested. The first *nfds* descriptors shall  
 31097 be checked in each set; that is, the descriptors from zero through *nfds*–1 in the descriptor sets  
 31098 shall be examined.

31099 If the *readfds* argument is not a null pointer, it points to an object of type **fd\_set** that on input  
 31100 specifies the file descriptors to be checked for being ready to read, and on output indicates  
 31101 which file descriptors are ready to read.

31102 If the *writefds* argument is not a null pointer, it points to an object of type **fd\_set** that on input  
 31103 specifies the file descriptors to be checked for being ready to write, and on output indicates  
 31104 which file descriptors are ready to write.

31105 If the *errorfds* argument is not a null pointer, it points to an object of type **fd\_set** that on input  
 31106 specifies the file descriptors to be checked for error conditions pending, and on output indicates  
 31107 which file descriptors have error conditions pending.

31108 Upon successful completion, the *pselect()* or *select()* function shall modify the objects pointed to  
 31109 by the *readfds*, *writefds*, and *errorfds* arguments to indicate which file descriptors are ready for  
 31110 reading, ready for writing, or have an error condition pending, respectively, and shall return the  
 31111 total number of ready descriptors in all the output sets. For each file descriptor less than *nfds*, the

31112 corresponding bit shall be set on successful completion if it was set on input and the associated  
31113 condition is true for that file descriptor.

31114 If none of the selected descriptors are ready for the requested operation, the *pselect()* or *select()* |  
31115 function shall block until at least one of the requested operations becomes ready, until the |  
31116 *timeout* occurs, or until interrupted by a signal. The *timeout* parameter controls how long the |  
31117 *pselect()* or *select()* function shall take before timing out. If the *timeout* parameter is not a null |  
31118 pointer, it specifies a maximum interval to wait for the selection to complete. If the specified  
31119 time interval expires without any requested operation becoming ready, the function shall return.  
31120 If the *timeout* parameter is a null pointer, then the call to *pselect()* or *select()* shall block  
31121 indefinitely until at least one descriptor meets the specified criteria. To effect a poll, the *timeout*  
31122 parameter should not be a null pointer, and should point to a zero-valued **timespec** structure.

31123 The use of a timeout does not affect any pending timers set up by *alarm()*, *ualarm()*, or  
31124 *setitimer()*.

31125 Implementations may place limitations on the maximum timeout interval supported. All  
31126 implementations shall support a maximum timeout interval of at least 31 days. If the *timeout*  
31127 argument specifies a timeout interval greater than the implementation-defined maximum value,  
31128 the maximum value shall be used as the actual timeout value. Implementations may also place  
31129 limitations on the granularity of timeout intervals. If the requested timeout interval requires a  
31130 finer granularity than the implementation supports, the actual timeout interval shall be rounded  
31131 up to the next supported value.

31132 If *sigmask* is not a null pointer, then the *pselect()* function shall replace the signal mask of the  
31133 process by the set of signals pointed to by *sigmask* before examining the descriptors, and shall  
31134 restore the signal mask of the process before returning.

31135 A descriptor shall be considered ready for reading when a call to an input function with  
31136 O\_NONBLOCK clear would not block, whether or not the function would transfer data  
31137 successfully. (The function might return data, an end-of-file indication, or an error other than  
31138 one indicating that it is blocked, and in each of these cases the descriptor shall be considered  
31139 ready for reading.)

31140 A descriptor shall be considered ready for writing when a call to an output function with  
31141 O\_NONBLOCK clear would not block, whether or not the function would transfer data  
31142 successfully.

31143 If a socket has a pending error, it shall be considered to have an exceptional condition pending.  
31144 Otherwise, what constitutes an exceptional condition is file type-specific. For a file descriptor for  
31145 use with a socket, it is protocol-specific except as noted below. For other file types it is  
31146 implementation-defined. If the operation is meaningless for a particular file type, *pselect()* or  
31147 *select()* shall indicate that the descriptor is ready for read or write operations, and shall indicate  
31148 that the descriptor has no exceptional condition pending.

31149 If a descriptor refers to a socket, the implied input function is the *recvmsg()* function with  
31150 parameters requesting normal and ancillary data, such that the presence of either type shall  
31151 cause the socket to be marked as readable. The presence of out-of-band data shall be checked if |  
31152 the socket option SO\_OOBINLINE has been enabled, as out-of-band data is enqueued with |  
31153 normal data. If the socket is currently listening, then it shall be marked as readable if an |  
31154 incoming connection request has been received, and a call to the *accept()* function shall complete |  
31155 without blocking. |

31156 If a descriptor refers to a socket, the implied output function is the *sendmsg()* function supplying  
31157 an amount of normal data equal to the current value of the SO\_SNDLOWAT option for the  
31158 socket. If a non-blocking call to the *connect()* function has been made for a socket, and the  
31159 connection attempt has either succeeded or failed leaving a pending error, the socket shall be

31160 marked as writable.

31161 A socket shall be considered to have an exceptional condition pending if a receive operation  
 31162 with `O_NONBLOCK` clear for the open file description and with the `MSG_OOB` flag set would  
 31163 return out-of-band data without blocking. (It is protocol-specific whether the `MSG_OOB` flag  
 31164 would be used to read out-of-band data.) A socket shall also be considered to have an  
 31165 exceptional condition pending if an out-of-band data mark is present in the receive queue. Other  
 31166 circumstances under which a socket may be considered to have an exceptional condition  
 31167 pending are protocol-specific and implementation-defined.

31168 If the `readfds`, `writefds`, and `errorfds` arguments are all null pointers and the `timeout` argument is not  
 31169 a null pointer, the `pselect()` or `select()` function shall block for the time specified, or until |  
 31170 interrupted by a signal. If the `readfds`, `writefds`, and `errorfds` arguments are all null pointers and the |  
 31171 `timeout` argument is a null pointer, the `pselect()` or `select()` function shall block until interrupted |  
 31172 by a signal. |

31173 File descriptors associated with regular files shall always select true for ready to read, ready to  
 31174 write, and error conditions.

31175 On failure, the objects pointed to by the `readfds`, `writefds`, and `errorfds` arguments shall not be |  
 31176 modified. If the timeout interval expires without the specified condition being true for any of the |  
 31177 specified file descriptors, the objects pointed to by the `readfds`, `writefds`, and `errorfds` arguments |  
 31178 shall have all bits set to 0. |

31179 File descriptor masks of type `fd_set` can be initialized and tested with `FD_CLR()`, `FD_ISSET()`,  
 31180 `FD_SET()`, and `FD_ZERO()`. It is unspecified whether each of these is a macro or a function. If a  
 31181 macro definition is suppressed in order to access an actual function, or a program defines an  
 31182 external identifier with any of these names, the behavior is undefined.

31183 `FD_CLR(fd, fdsetp)` shall remove the file descriptor `fd` from the set pointed to by `fdsetp`. If `fd` is not |  
 31184 a member of this set, there shall be no effect on the set, nor will an error be returned.

31185 `FD_ISSET(fd, fdsetp)` shall evaluate to non-zero if the file descriptor `fd` is a member of the set |  
 31186 pointed to by `fdsetp`, and shall evaluate to zero otherwise.

31187 `FD_SET(fd, fdsetp)` shall add the file descriptor `fd` to the set pointed to by `fdsetp`. If the file |  
 31188 descriptor `fd` is already in this set, there shall be no effect on the set, nor will an error be returned.

31189 `FD_ZERO(fdsetp)` shall initialize the descriptor set pointed to by `fdsetp` to the null set. No error is |  
 31190 returned if the set is not empty at the time `FD_ZERO()` is invoked.

31191 The behavior of these macros is undefined if the `fd` argument is less than 0 or greater than or  
 31192 equal to `FD_SETSIZE`, or if `fd` is not a valid file descriptor, or if any of the arguments are  
 31193 expressions with side effects.

#### 31194 RETURN VALUE

31195 Upon successful completion, the `pselect()` and `select()` functions shall return the total number of  
 31196 bits set in the bit masks. Otherwise, `-1` shall be returned, and shall set `errno` to indicate the error.

31197 `FD_CLR()`, `FD_SET()`, and `FD_ZERO()` not return a value. `FD_ISSET()` shall return a non-zero  
 31198 value if the bit for the file descriptor `fd` is set in the file descriptor set pointed to by `fdset`, and 0  
 31199 otherwise.

#### 31200 ERRORS

31201 Under the following conditions, `pselect()` and `select()` shall fail and set `errno` to:

31202 [EBADF] One or more of the file descriptor sets specified a file descriptor that is not a  
 31203 valid open file descriptor.

- 31204 [EINTR] The function was interrupted before any of the selected events occurred and  
31205 before the timeout interval expired.
- 31206 XSI If SA\_RESTART has been set for the interrupting signal, it is implementation-  
31207 defined whether the function restarts or returns with [EINTR].
- 31208 [EINVAL] An invalid timeout interval was specified.
- 31209 [EINVAL] The *nfds* argument is less than 0 or greater than FD\_SETSIZE.
- 31210 XSR [EINVAL] One of the specified file descriptors refers to a STREAM or multiplexer that is  
31211 linked (directly or indirectly) downstream from a multiplexer.

31212 **EXAMPLES**

31213 None.

31214 **APPLICATION USAGE**

31215 None.

31216 **RATIONALE**

31217 In previous versions of the Single UNIX Specification, the *select()* function was defined in the  
31218 `<sys/time.h>` header. This is now changed to `<sys/select.h>`. The rationale for this change was  
31219 as follows: the introduction of the *pselect()* function included the `<sys/select.h>` header and the  
31220 `<sys/select.h>` header defines all the related definitions for the *pselect()* and *select()* functions.  
31221 Backwards-compatibility to existing XSI implementations is handled by allowing `<sys/time.h>`  
31222 to include `<sys/select.h>`.

31223 **FUTURE DIRECTIONS**

31224 None.

31225 **SEE ALSO**

31226 *accept()*, *alarm()*, *connect()*, *fcntl()*, *poll()*, *read()*, *recvmsg()*, *sendmsg()*, *setitimer()*, *ualarm()*,  
31227 *write()*, the Base Definitions volume of IEEE Std 1003.1-200x, `<sys/select.h>`, `<sys/time.h>`

31228 **CHANGE HISTORY**

31229 First released in Issue 4, Version 2.

31230 **Issue 5**

31231 Moved from X/OPEN UNIX extension to BASE.

31232 In the ERRORS section, the text has been changed to indicate that [EINVAL] is returned when  
31233 *nfds* is less than 0 or greater than FD\_SETSIZE. It previously stated less than 0, or greater than or  
31234 equal to FD\_SETSIZE.

31235 Text about *timeout* is moved from the APPLICATION USAGE section to the DESCRIPTION.31236 **Issue 6**31237 The Open Group Corrigendum U026/6 is applied, changing the occurrences of *readfs* and *writfs*  
31238 in the *select()* DESCRIPTION to be *readfds* and *writefds*.

31239 Text referring to sockets is added to the DESCRIPTION.

31240 The DESCRIPTION and ERRORS sections are updated so that references to STREAMS are  
31241 marked as part of the XSI STREAMS Option Group.

31242 The following new requirements on POSIX implementations derive from alignment with the  
31243 Single UNIX Specification:

- 31244 • These functions are now mandatory.

31245 The *pselect()* function is added for alignment with IEEE Std 1003.1g-2000 and additional detail  
31246 related to sockets semantics is added to the DESCRIPTION.



- 31247        The *select()* function now requires inclusion of `<sys/select.h>`.
- 31248        The **restrict** keyword is added to the *select()* prototype for alignment with the
- 31249        ISO/IEC 9899:1999 standard.

31250 **NAME**

31251 pthread\_atfork — register fork handlers

31252 **SYNOPSIS**

31253 THR #include &lt;pthread.h&gt;

31254 int pthread\_atfork(void (\*prepare)(void), void (\*parent)(void),  
31255 void (\*child)(void));

31256

31257 **DESCRIPTION**

31258 The *pthread\_atfork()* function shall declare fork handlers to be called before and after *fork()*, in  
31259 the context of the thread that called *fork()*. The *prepare* fork handler shall be called before *fork()*  
31260 processing commences. The *parent* fork handle shall be called after *fork()* processing completes  
31261 in the parent process. The *child* fork handler shall be called after *fork()* processing completes in  
31262 the child process. If no handling is desired at one or more of these three points, the  
31263 corresponding fork handler address(es) may be set to NULL.

31264 The order of calls to *pthread\_atfork()* is significant. The *parent* and *child* fork handlers shall be  
31265 called in the order in which they were established by calls to *pthread\_atfork()*. The *prepare* fork  
31266 handlers shall be called in the opposite order.

31267 **RETURN VALUE**

31268 Upon successful completion, *pthread\_atfork()* shall return a value of zero; otherwise, an error  
31269 number shall be returned to indicate the error.

31270 **ERRORS**31271 The *pthread\_atfork()* function shall fail if:

31272 [ENOMEM] Insufficient table space exists to record the fork handler addresses.

31273 The *pthread\_atfork()* function shall not return an error code of [EINTR].31274 **EXAMPLES**

31275 None.

31276 **APPLICATION USAGE**

31277 None.

31278 **RATIONALE**

31279 There are at least two serious problems with the semantics of *fork()* in a multi-threaded  
31280 program. One problem has to do with state (for example, memory) covered by mutexes.  
31281 Consider the case where one thread has a mutex locked and the state covered by that mutex is  
31282 inconsistent while another thread calls *fork()*. In the child, the mutex is in the locked state  
31283 (locked by a nonexistent thread and thus can never be unlocked). Having the child simply  
31284 reinitialize the mutex is unsatisfactory since this approach does not resolve the question about  
31285 how to correct or otherwise deal with the inconsistent state in the child.

31286 It is suggested that programs that use *fork()* call an *exec* function very soon afterwards in the  
31287 child process, thus resetting all states. In the meantime, only a short list of async-signal-safe  
31288 library routines are promised to be available.

31289 Unfortunately, this solution does not address the needs of multi-threaded libraries. Application  
31290 programs may not be aware that a multi-threaded library is in use, and they feel free to call any  
31291 number of library routines between the *fork()* and *exec* calls, just as they always have. Indeed,  
31292 they may be extant single-threaded programs and cannot, therefore, be expected to obey new  
31293 restrictions imposed by the threads library.

31294 On the other hand, the multi-threaded library needs a way to protect its internal state during  
 31295 *fork()* in case it is re-entered later in the child process. The problem arises especially in multi-  
 31296 threaded I/O libraries, which are almost sure to be invoked between the *fork()* and *exec* calls to  
 31297 effect I/O redirection. The solution may require locking mutex variables during *fork()*, or it may  
 31298 entail simply resetting the state in the child after the *fork()* processing completes.

31299 The *pthread\_atfork()* function provides multi-threaded libraries with a means to protect  
 31300 themselves from innocent application programs that call *fork()*, and it provides multi-threaded  
 31301 application programs with a standard mechanism for protecting themselves from *fork()* calls in  
 31302 a library routine or the application itself.

31303 The expected usage is that the *prepare* handler acquires all mutex locks and the other two fork  
 31304 handlers release them.

31305 For example, an application can supply a *prepare* routine that acquires the necessary mutexes the  
 31306 library maintains and supply *child* and *parent* routines that release those mutexes, thus ensuring  
 31307 that the child gets a consistent snapshot of the state of the library (and that no mutexes are left  
 31308 stranded). Alternatively, some libraries might be able to supply just a *child* routine that  
 31309 reinitializes the mutexes in the library and all associated states to some known value (for  
 31310 example, what it was when the image was originally executed).

31311 When *fork()* is called, only the calling thread is duplicated in the child process. Synchronization  
 31312 variables remain in the same state in the child as they were in the parent at the time *fork()* was  
 31313 called. Thus, for example, mutex locks may be held by threads that no longer exist in the child  
 31314 process, and any associated states may be inconsistent. The parent process may avoid this by  
 31315 explicit code that acquires and releases locks critical to the child via *pthread\_atfork()*. In addition,  
 31316 any critical threads need to be recreated and reinitialized to the proper state in the child (also via  
 31317 *pthread\_atfork()*).

31318 A higher-level package may acquire locks on its own data structures before invoking lower-level  
 31319 packages. Under this scenario, the order specified for fork handler calls allows a simple rule of  
 31320 initialization for avoiding package deadlock: a package initializes all packages on which it  
 31321 depends before it calls the *pthread\_atfork()* function for itself.

#### 31322 **FUTURE DIRECTIONS**

31323 None.

#### 31324 **SEE ALSO**

31325 *atexit()*, *fork()*, the Base Definitions volume of IEEE Std 1003.1-200x, <sys/types.h>

#### 31326 **CHANGE HISTORY**

31327 First released in Issue 5. Derived from the POSIX Threads Extension.

31328 IEEE PASC Interpretation 1003.1c #4 is applied.

#### 31329 **Issue 6**

31330 The *pthread\_atfork()* function is marked as part of the Threads option.

31331 The <pthread.h> header is added to the SYNOPSIS.

## 31332 NAME

31333 pthread\_attr\_destroy, pthread\_attr\_init — destroy and initialize threads attributes object

## 31334 SYNOPSIS

```
31335 THR #include <pthread.h>
```

```
31336 int pthread_attr_destroy(pthread_attr_t *attr);
```

```
31337 int pthread_attr_init(pthread_attr_t *attr);
```

31338

## 31339 DESCRIPTION

31340 The *pthread\_attr\_destroy()* function shall destroy a thread attributes object. An implementation |  
31341 may cause *pthread\_attr\_destroy()* to set *attr* to an implementation-defined invalid value. A |  
31342 destroyed *attr* attributes object can be reinitialized using *pthread\_attr\_init()*; the results of |  
31343 otherwise referencing the object after it has been destroyed are undefined. |

31344 The *pthread\_attr\_init()* function shall initialize a thread attributes object *attr* with the default |  
31345 value for all of the individual attributes used by a given implementation.

31346 The resulting attributes object (possibly modified by setting individual attribute values), when |  
31347 used by *pthread\_create()* defines the attributes of the thread created. A single attributes object can |  
31348 be used in multiple simultaneous calls to *pthread\_create()*. Results are undefined if |  
31349 *pthread\_attr\_init()* is called specifying an already initialized *attr* attributes object. |

## 31350 RETURN VALUE

31351 Upon successful completion, *pthread\_attr\_destroy()* and *pthread\_attr\_init()* shall return a value of |  
31352 0; otherwise, an error number shall be returned to indicate the error.

## 31353 ERRORS

31354 The *pthread\_attr\_init()* function shall fail if:

31355 [ENOMEM] Insufficient memory exists to initialize the thread attributes object.

31356 These functions shall not return an error code of [EINTR].

## 31357 EXAMPLES

31358 None.

## 31359 APPLICATION USAGE

31360 None.

## 31361 RATIONALE

31362 Attributes objects are provided for threads, mutexes, and condition variables as a mechanism to |  
31363 support probable future standardization in these areas without requiring that the function itself |  
31364 be changed.

31365 Attributes objects provide clean isolation of the configurable aspects of threads. For example, |  
31366 “stack size” is an important attribute of a thread, but it cannot be expressed portably. When |  
31367 porting a threaded program, stack sizes often need to be adjusted. The use of attributes objects |  
31368 can help by allowing the changes to be isolated in a single place, rather than being spread across |  
31369 every instance of thread creation.

31370 Attributes objects can be used to set up “classes” of threads with similar attributes; for example, |  
31371 “threads with large stacks and high priority” or “threads with minimal stacks”. These classes |  
31372 can be defined in a single place and then referenced wherever threads need to be created. |  
31373 Changes to “class” decisions become straightforward, and detailed analysis of each |  
31374 *pthread\_create()* call is not required.

31375 The attributes objects are defined as opaque types as an aid to extensibility. If these objects had |  
31376 been specified as structures, adding new attributes would force recompilation of all multi-

31377 threaded programs when the attributes objects are extended; this might not be possible if  
31378 different program components were supplied by different vendors.

31379 Additionally, opaque attributes objects present opportunities for improving performance.  
31380 Argument validity can be checked once when attributes are set, rather than each time a thread is  
31381 created. Implementations often need to cache kernel objects that are expensive to create.  
31382 Opaque attributes objects provide an efficient mechanism to detect when cached objects become  
31383 invalid due to attribute changes.

31384 Since assignment is not necessarily defined on a given opaque type, implementation-defined  
31385 default values cannot be defined in a portable way. The solution to this problem is to allow  
31386 attributes objects to be initialized dynamically by attributes object initialization functions, so  
31387 that default values can be supplied automatically by the implementation.

31388 The following proposal was provided as a suggested alternative to the supplied attributes:

- 31389 1. Maintain the style of passing a parameter formed by the bitwise-inclusive OR of flags to  
31390 the initialization routines (*pthread\_create()*, *pthread\_mutex\_init()*, *pthread\_cond\_init()*). The  
31391 parameter containing the flags should be an opaque type for extensibility. If no flags are  
31392 set in the parameter, then the objects are created with default characteristics. An  
31393 implementation may specify implementation-defined flag values and associated behavior.
- 31394 2. If further specialization of mutexes and condition variables is necessary, implementations  
31395 may specify additional procedures that operate on the **pthread\_mutex\_t** and  
31396 **pthread\_cond\_t** objects (instead of on attributes objects).

31397 The difficulties with this solution are:

- 31398 1. A bitmask is not opaque if bits have to be set into bitvector attributes objects using  
31399 explicitly-coded bitwise-inclusive OR operations. If the set of options exceeds an **int**,  
31400 application programmers need to know the location of each bit. If bits are set or read by  
31401 encapsulation (that is, *get* and *set* functions), then the bitmask is merely an  
31402 implementation of attributes objects as currently defined and should not be exposed to the  
31403 programmer.
- 31404 2. Many attributes are not Boolean or very small integral values. For example, scheduling  
31405 policy may be placed in 3-bit or 4-bit, but priority requires 5-bit or more, thereby taking up  
31406 at least 8 bits out of a possible 16 bits on machines with 16-bit integers. Because of this, the  
31407 bitmask can only reasonably control whether particular attributes are set or not, and it  
31408 cannot serve as the repository of the value itself. The value needs to be specified as a  
31409 function parameter (which is non-extensible), or by setting a structure field (which is non-  
31410 opaque), or by *get* and *set* functions (making the bitmask a redundant addition to the  
31411 attributes objects).

31412 Stack size is defined as an optional attribute because the very notion of a stack is inherently  
31413 machine-dependent. Some implementations may not be able to change the size of the stack, for  
31414 example, and others may not need to because stack pages may be discontinuous and can be  
31415 allocated and released on demand.

31416 The attribute mechanism has been designed in large measure for extensibility. Future extensions  
31417 to the attribute mechanism or to any attributes object defined in this volume of  
31418 IEEE Std 1003.1-200x has to be done with care so as not to affect binary-compatibility.

31419 Attributes objects, even if allocated by means of dynamic allocation functions such as *malloc()*,  
31420 may have their size fixed at compile time. This means, for example, a *pthread\_create()* in an  
31421 implementation with extensions to the **pthread\_attr\_t** cannot look beyond the area that the  
31422 binary application assumes is valid. This suggests that implementations should maintain a size  
31423 field in the attributes object, as well as possibly version information, if extensions in different

31424 directions (possibly by different vendors) are to be accommodated.

## 31425 FUTURE DIRECTIONS

31426 None.

## 31427 SEE ALSO

31428 *pthread\_attr\_getstackaddr()*, *pthread\_attr\_getstacksize()*, *pthread\_attr\_getdetachstate()*,  
31429 *pthread\_create()*, the Base Definitions volume of IEEE Std 1003.1-200x, <pthread.h>

## 31430 CHANGE HISTORY

31431 First released in Issue 5. Included for alignment with the POSIX Threads Extension.

## 31432 Issue 6

31433 The *pthread\_attr\_destroy()* and *pthread\_attr\_init()* functions marked as part of the Threads  
31434 option.

31435 IEEE PASC Interpretation 1003.1 #107 is applied, noting that the effect of initializing an already  
31436 initialized thread attributes object is undefined.

31437 **NAME**

31438 pthread\_attr\_getdetachstate, pthread\_attr\_setdetachstate — get and set detachstate attribute

31439 **SYNOPSIS**

31440 THR #include <pthread.h>

31441 int pthread\_attr\_getdetachstate(const pthread\_attr\_t \*attr,

31442 int \*detachstate);

31443 int pthread\_attr\_setdetachstate(pthread\_attr\_t \*attr, int detachstate);

31444

31445 **DESCRIPTION**

31446 The *detachstate* attribute controls whether the thread is created in a detached state. If the thread  
31447 is created detached, then use of the ID of the newly created thread by the *pthread\_detach()* or  
31448 *pthread\_join()* function is an error.

31449 The *pthread\_attr\_getdetachstate()* and *pthread\_attr\_setdetachstate()* functions, respectively, shall  
31450 get and set the *detachstate* attribute in the *attr* object.

31451 For *pthread\_attr\_getdetachstate()*, *detachstate* shall be set to either  
31452 PTHREAD\_CREATE\_DETACHED or PTHREAD\_CREATE\_JOINABLE.

31453 For *pthread\_attr\_setdetachstate()*, the application shall set *detachstate* to either  
31454 PTHREAD\_CREATE\_DETACHED or PTHREAD\_CREATE\_JOINABLE.

31455 A value of PTHREAD\_CREATE\_DETACHED shall cause all threads created with *attr* to be in  
31456 the detached state, whereas using a value of PTHREAD\_CREATE\_JOINABLE shall cause all  
31457 threads created with *attr* to be in the joinable state. The default value of the *detachstate* attribute  
31458 shall be PTHREAD\_CREATE\_JOINABLE.

31459 **RETURN VALUE**

31460 Upon successful completion, *pthread\_attr\_getdetachstate()* and *pthread\_attr\_setdetachstate()* shall  
31461 return a value of 0; otherwise, an error number shall be returned to indicate the error.

31462 The *pthread\_attr\_getdetachstate()* function stores the value of the *detachstate* attribute in *detachstate*  
31463 if successful.

31464 **ERRORS**

31465 The *pthread\_attr\_setdetachstate()* function shall fail if:

31466 [EINVAL] The value of *detachstate* was not valid

31467 These functions shall not return an error code of [EINTR].

31468 **EXAMPLES**

31469 None.

31470 **APPLICATION USAGE**

31471 None.

31472 **RATIONALE**

31473 None.

31474 **FUTURE DIRECTIONS**

31475 None.

31476 **SEE ALSO**

31477 *pthread\_attr\_destroy()*, *pthread\_attr\_getstackaddr()*, *pthread\_attr\_getstacksize()*, *pthread\_create()*, the  
31478 Base Definitions volume of IEEE Std 1003.1-200x, <pthread.h>

31479 **CHANGE HISTORY**

31480 First released in Issue 5. Included for alignment with the POSIX Threads Extension.

31481 **Issue 6**

31482 The *pthread\_attr\_setdetachstate()* and *pthread\_attr\_getdetachstate()* functions are marked as part of  
31483 the Threads option.

31484 The DESCRIPTION is updated to avoid use of the term “must” for application requirements.



31485 **NAME**

31486 pthread\_attr\_getguardsize, pthread\_attr\_setguardsize — get and set the thread guardsize  
 31487 attribute

31488 **SYNOPSIS**

```
31489 XSI #include <pthread.h>
31490
31490 int pthread_attr_getguardsize(const pthread_attr_t *restrict attr,
31491                             size_t *restrict guardsize);
31492 int pthread_attr_setguardsize(pthread_attr_t *attr,
31493                             size_t guardsize);
31494
```

31495 **DESCRIPTION**

31496 The *pthread\_attr\_getguardsize()* function shall get the *guardsize* attribute in the *attr* object. This  
 31497 attribute shall be returned in the *guardsize* parameter.

31498 The *pthread\_attr\_setguardsize()* function shall set the *guardsize* attribute in the *attr* object. The new  
 31499 value of this attribute shall be obtained from the *guardsize* parameter. If *guardsize* is zero, a guard  
 31500 area shall not be provided for threads created with *attr*. If *guardsize* is greater than zero, a guard  
 31501 area of at least size *guardsize* bytes shall be provided for each thread created with *attr*.

31502 The *guardsize* attribute controls the size of the guard area for the created thread's stack. The  
 31503 *guardsize* attribute provides protection against overflow of the stack pointer. If a thread's stack is  
 31504 created with guard protection, the implementation allocates extra memory at the overflow end  
 31505 of the stack as a buffer against stack overflow of the stack pointer. If an application overflows  
 31506 into this buffer an error shall result (possibly in a SIGSEGV signal being delivered to the thread).

31507 A conforming implementation may round up the value contained in *guardsize* to a multiple of  
 31508 the configurable system variable {PAGESIZE} (see <sys/mman.h>). If an implementation  
 31509 rounds up the value of *guardsize* to a multiple of {PAGESIZE}, a call to *pthread\_attr\_getguardsize()*  
 31510 specifying *attr* shall store in the *guardsize* parameter the guard size specified by the previous  
 31511 *pthread\_attr\_setguardsize()* function call.

31512 The default value of the *guardsize* attribute is {PAGESIZE} bytes. The actual value of {PAGESIZE}  
 31513 is implementation-defined.

31514 If the *stackaddr* or *stack* attribute has been set (that is, the caller is allocating and managing its  
 31515 own thread stacks), the *guardsize* attribute shall be ignored and no protection shall be provided  
 31516 by the implementation. It is the responsibility of the application to manage stack overflow along  
 31517 with stack allocation and management in this case.

31518 **RETURN VALUE**

31519 If successful, the *pthread\_attr\_getguardsize()* and *pthread\_attr\_setguardsize()* functions shall return  
 31520 zero; otherwise, an error number shall be returned to indicate the error.

31521 **ERRORS**

31522 The *pthread\_attr\_getguardsize()* and *pthread\_attr\_setguardsize()* functions shall fail if:

31523 [EINVAL] The attribute *attr* is invalid.

31524 [EINVAL] The parameter *guardsize* is invalid.

31525 These functions shall not return an error code of [EINTR].

## 31526 EXAMPLES

31527 None.

## 31528 APPLICATION USAGE

31529 None.

## 31530 RATIONALE

31531 The *guardsize* attribute is provided to the application for two reasons:

- 31532 1. Overflow protection can potentially result in wasted system resources. An application  
31533 that creates a large number of threads, and which knows its threads never overflow their  
31534 stack, can save system resources by turning off guard areas.
- 31535 2. When threads allocate large data structures on the stack, large guard areas may be needed  
31536 to detect stack overflow.

## 31537 FUTURE DIRECTIONS

31538 None.

## 31539 SEE ALSO

31540 The Base Definitions volume of IEEE Std 1003.1-200x, `<pthread.h>`, `<sys/mman.h>`

## 31541 CHANGE HISTORY

31542 First released in Issue 5.

### 31543 Issue 6

31544 In the ERRORS section, a third [EINVAL] error condition is removed as it is covered by the  
31545 second error condition.

31546 The **restrict** keyword is added to the *pthread\_attr\_getguardsize()* prototype for alignment with the  
31547 ISO/IEC 9899:1999 standard.

31548 **NAME**

31549 pthread\_attr\_getinheritsched, pthread\_attr\_setinheritsched — get and set inheritsched attribute  
 31550 (**REALTIME THREADS**)

31551 **SYNOPSIS**

31552 THR TPS #include <pthread.h>

```
31553 int pthread_attr_getinheritsched(const pthread_attr_t *restrict attr,
31554                               int *restrict inheritsched);
31555 int pthread_attr_setinheritsched(pthread_attr_t *attr,
31556                               int inheritsched);
31557
```

31558 **DESCRIPTION**

31559 The *pthread\_attr\_getinheritsched()*, and *pthread\_attr\_setinheritsched()* functions, respectively, shall  
 31560 get and set the *inheritsched* attribute in the *attr* argument.

31561 When the attributes objects are used by *pthread\_create()*, the *inheritsched* attribute determines  
 31562 how the other scheduling attributes of the created thread shall be set.

31563 **PTHREAD\_INHERIT\_SCHED**

31564 Specifies that the thread scheduling attributes shall be inherited from the creating thread,  
 31565 and the scheduling attributes in this *attr* argument shall be ignored.

31566 **PTHREAD\_EXPLICIT\_SCHED**

31567 Specifies that the thread scheduling attributes shall be set to the corresponding values from  
 31568 this attributes object.

31569 The symbols **PTHREAD\_INHERIT\_SCHED** and **PTHREAD\_EXPLICIT\_SCHED** are defined in  
 31570 the <**pthread.h**> header.

31571 The following “thread scheduling attributes” defined by IEEE Std 1003.1-200x are affected by  
 31572 the *inheritsched* attribute: scheduling policy (*schedpolicy*), scheduling parameters (*schedparam*),  
 31573 and scheduling contention scope (*contentionscope*).

31574 **RETURN VALUE**

31575 If successful, the *pthread\_attr\_getinheritsched()* and *pthread\_attr\_setinheritsched()* functions shall  
 31576 return zero; otherwise, an error number shall be returned to indicate the error.

31577 **ERRORS**

31578 The *pthread\_attr\_setinheritsched()* function may fail if:

31579 [EINVAL] The value of *inheritsched* is not valid.

31580 [ENOTSUP] An attempt was made to set the attribute to an unsupported value.

31581 These functions shall not return an error code of [EINTR].

31582 **EXAMPLES**

31583 None.

31584 **APPLICATION USAGE**

31585 After these attributes have been set, a thread can be created with the specified attributes using  
 31586 *pthread\_create()*. Using these routines does not affect the current running thread.

31587 **RATIONALE**

31588 None.

## 31589 FUTURE DIRECTIONS

31590 None.

## 31591 SEE ALSO

31592 *pthread\_attr\_destroy()*, *pthread\_attr\_getscope()*, *pthread\_attr\_getschedpolicy()*,  
31593 *pthread\_attr\_getschedparam()*, *pthread\_create()*, the Base Definitions volume of  
31594 IEEE Std 1003.1-200x, <pthread.h>, <sched.h>

## 31595 CHANGE HISTORY

31596 First released in Issue 5. Included for alignment with the POSIX Threads Extension.

31597 Marked as part of the Realtime Threads Feature Group.

## 31598 Issue 6

31599 The *pthread\_attr\_getinheritsched()* and *pthread\_attr\_setinheritsched()* functions are marked as part  
31600 of the Threads and Thread Execution Scheduling options.

31601 The [ENOSYS] error condition has been removed as stubs need not be provided if an  
31602 implementation does not support the Thread Execution Scheduling option.

31603 The **restrict** keyword is added to the *pthread\_attr\_getinheritsched()* prototype for alignment with  
31604 the ISO/IEC 9899:1999 standard.

31605 **NAME**

31606 pthread\_attr\_getschedparam, pthread\_attr\_setschedparam — get and set schedparam attribute

31607 **SYNOPSIS**

31608 THR #include <pthread.h>

```
31609 int pthread_attr_getschedparam(const pthread_attr_t *restrict attr,
31610 struct sched_param *restrict param);
```

```
31611 int pthread_attr_setschedparam(pthread_attr_t *restrict attr,
31612 const struct sched_param *restrict param);
```

31613

31614 **DESCRIPTION**

31615 The *pthread\_attr\_getschedparam()*, and *pthread\_attr\_setschedparam()* functions, respectively, shall  
 31616 get and set the scheduling parameter attributes in the *attr* argument. The contents of the *param* |  
 31617 structure are defined in the <sched.h> header. For the SCHED\_FIFO and SCHED\_RR policies, |  
 31618 the only required member of *param* is *sched\_priority*.

31619 TSP For the SCHED\_SPORADIC policy, the required members of the *param* structure are  
 31620 *sched\_priority*, *sched\_ss\_low\_priority*, *sched\_ss\_repl\_period*, *sched\_ss\_init\_budget*, and  
 31621 *sched\_ss\_max\_repl*. The specified *sched\_ss\_repl\_period* must be greater than or equal to the  
 31622 specified *sched\_ss\_init\_budget* for the function to succeed; if it is not, then the function shall fail.  
 31623 The value of *sched\_ss\_max\_repl* shall be within the inclusive range [1,{SS\_REPL\_MAX}] for the  
 31624 function to succeed; if not, the function shall fail.

31625 **RETURN VALUE**

31626 If successful, the *pthread\_attr\_getschedparam()* and *pthread\_attr\_setschedparam()* functions shall  
 31627 return zero; otherwise, an error number shall be returned to indicate the error.

31628 **ERRORS**

31629 The *pthread\_attr\_setschedparam()* function may fail if:

31630 [EINVAL] The value of *param* is not valid.

31631 [ENOTSUP] An attempt was made to set the attribute to an unsupported value.

31632 These functions shall not return an error code of [EINTR].

31633 **EXAMPLES**

31634 None.

31635 **APPLICATION USAGE**

31636 After these attributes have been set, a thread can be created with the specified attributes using  
 31637 *pthread\_create()*. Using these routines does not affect the current running thread.

31638 **RATIONALE**

31639 None.

31640 **FUTURE DIRECTIONS**

31641 None.

31642 **SEE ALSO**

31643 *pthread\_attr\_destroy()*, *pthread\_attr\_getscope()*, *pthread\_attr\_getinheritsched()*,  
 31644 *pthread\_attr\_getschedpolicy()*, *pthread\_create()*, the Base Definitions volume of  
 31645 IEEE Std 1003.1-200x, <pthread.h>, <sched.h>

## 31646 CHANGE HISTORY

31647 First released in Issue 5. Included for alignment with the POSIX Threads Extension.

## 31648 Issue 6

31649 The *pthread\_attr\_getschedparam()* and *pthread\_attr\_setschedparam()* functions are marked as part  
31650 of the Threads option.

31651 The SCHED\_SPORADIC scheduling policy is added for alignment with IEEE Std 1003.1d-1999.

31652 The **restrict** keyword is added to the *pthread\_attr\_getschedparam()* and  
31653 *pthread\_attr\_setschedparam()* prototypes for alignment with the ISO/IEC 9899:1999 standard.

31654 **NAME**

31655 pthread\_attr\_getschedpolicy, pthread\_attr\_setschedpolicy — get and set schedpolicy attribute  
 31656 (**REALTIME THREADS**)

31657 **SYNOPSIS**

```
31658 THR TPS #include <pthread.h>
31659
31659 int pthread_attr_getschedpolicy(const pthread_attr_t *restrict attr,
31660 int *restrict policy);
31661 int pthread_attr_setschedpolicy(pthread_attr_t *attr, int policy);
31662
```

31663 **DESCRIPTION**

31664 The *pthread\_attr\_getschedpolicy()* and *pthread\_attr\_setschedpolicy()* functions, respectively, shall  
 31665 get and set the *schedpolicy* attribute in the *attr* argument.

31666 The supported values of *policy* shall include SCHED\_FIFO, SCHED\_RR, and SCHED\_OTHER, |  
 31667 which are defined in the <*sched.h*> header. When threads executing with the scheduling policy |  
 31668 TSP SCHED\_FIFO, SCHED\_RR, or SCHED\_SPORADIC are waiting on a mutex, they shall acquire |  
 31669 the mutex in priority order when the mutex is unlocked. |

31670 **RETURN VALUE**

31671 If successful, the *pthread\_attr\_getschedpolicy()* and *pthread\_attr\_setschedpolicy()* functions shall  
 31672 return zero; otherwise, an error number shall be returned to indicate the error.

31673 **ERRORS**

31674 The *pthread\_attr\_setschedpolicy()* function may fail if:

- 31675 [EINVAL] The value of *policy* is not valid.
- 31676 [ENOTSUP] An attempt was made to set the attribute to an unsupported value.
- 31677 These functions shall not return an error code of [EINTR].

31678 **EXAMPLES**

31679 None.

31680 **APPLICATION USAGE**

31681 After these attributes have been set, a thread can be created with the specified attributes using  
 31682 *pthread\_create()*. Using these routines does not affect the current running thread.

31683 **RATIONALE**

31684 None.

31685 **FUTURE DIRECTIONS**

31686 None.

31687 **SEE ALSO**

31688 *pthread\_attr\_destroy()*, *pthread\_attr\_getscope()*, *pthread\_attr\_getinheritsched()*,  
 31689 *pthread\_attr\_getschedparam()*, *pthread\_create()*, the Base Definitions volume of  
 31690 IEEE Std 1003.1-200x, <*pthread.h*>, <*sched.h*>

31691 **CHANGE HISTORY**

- 31692 First released in Issue 5. Included for alignment with the POSIX Threads Extension.
- 31693 Marked as part of the Realtime Threads Feature Group.

31694 **Issue 6**

31695 The *pthread\_attr\_getschedpolicy()* and *pthread\_attr\_setschedpolicy()* functions are marked as part of  
31696 the Threads and Thread Execution Scheduling options.

31697 The [ENOSYS] error condition has been removed as stubs need not be provided if an  
31698 implementation does not support the Thread Execution Scheduling option.

31699 The SCHED\_SPORADIC scheduling policy is added for alignment with IEEE Std 1003.1d-1999.

31700 The **restrict** keyword is added to the *pthread\_attr\_getschedpolicy()* prototype for alignment with  
31701 the ISO/IEC 9899:1999 standard.



31702 **NAME**

31703 pthread\_attr\_getscope, pthread\_attr\_setscope — get and set contention scope attribute  
 31704 (**REALTIME THREADS**)

31705 **SYNOPSIS**

31706 THR TPS #include <pthread.h>

```
31707 int pthread_attr_getscope(const pthread_attr_t *restrict attr,
31708                          int *restrict contention_scope);
31709 int pthread_attr_setscope(pthread_attr_t *attr, int contention_scope);
31710
```

31711 **DESCRIPTION**

31712 The *pthread\_attr\_getscope()* and *pthread\_attr\_setscope()* functions, respectively, shall get and set  
 31713 the *contention\_scope* attribute in the *attr* object.

31714 The *contention\_scope* attribute may have the values PTHREAD\_SCOPE\_SYSTEM, signifying  
 31715 system scheduling contention scope, or PTHREAD\_SCOPE\_PROCESS, signifying process  
 31716 scheduling contention scope. The symbols PTHREAD\_SCOPE\_SYSTEM and  
 31717 PTHREAD\_SCOPE\_PROCESS are defined in the <pthread.h> header.

31718 **RETURN VALUE**

31719 If successful, the *pthread\_attr\_getscope()* and *pthread\_attr\_setscope()* functions shall return zero;  
 31720 otherwise, an error number shall be returned to indicate the error.

31721 **ERRORS**

31722 The *pthread\_attr\_setscope()* function may fail if:

31723 [EINVAL] The value of *contention\_scope* is not valid.

31724 [ENOTSUP] An attempt was made to set the attribute to an unsupported value.

31725 These functions shall not return an error code of [EINTR].

31726 **EXAMPLES**

31727 None.

31728 **APPLICATION USAGE**

31729 After these attributes have been set, a thread can be created with the specified attributes using  
 31730 *pthread\_create()*. Using these routines does not affect the current running thread.

31731 **RATIONALE**

31732 None.

31733 **FUTURE DIRECTIONS**

31734 None.

31735 **SEE ALSO**

31736 *pthread\_attr\_destroy()*, *pthread\_attr\_getinheritsched()*, *pthread\_attr\_getschedpolicy()*,  
 31737 *pthread\_attr\_getschedparam()*, *pthread\_create()*, the Base Definitions volume of  
 31738 IEEE Std 1003.1-200x, <pthread.h>, <sched.h>

31739 **CHANGE HISTORY**

31740 First released in Issue 5. Included for alignment with the POSIX Threads Extension.

31741 Marked as part of the Realtime Threads Feature Group.

31742 **Issue 6**

31743 The *pthread\_attr\_getscope()* and *pthread\_attr\_setscope()* functions are marked as part of the  
31744 Threads and Thread Execution Scheduling options.

31745 The [ENOSYS] error condition has been removed as stubs need not be provided if an  
31746 implementation does not support the Thread Execution Scheduling option.

31747 The **restrict** keyword is added to the *pthread\_attr\_getscope()* prototype for alignment with the  
31748 ISO/IEC 9899:1999 standard.

31749 **NAME**

31750 pthread\_attr\_getstack, pthread\_attr\_setstack — get and set stack attributes

31751 **SYNOPSIS**

31752 THR #include &lt;pthread.h&gt;

```

31753 TSA TSS int pthread_attr_getstack(const pthread_attr_t *restrict attr,
31754 void **restrict stackaddr, size_t *restrict stacksize);
31755 int pthread_attr_setstack(pthread_attr_t *attr, void *stackaddr,
31756 size_t stacksize);
31757

```

31758 **DESCRIPTION**

31759 The *pthread\_attr\_getstack()* and *pthread\_attr\_setstack()* functions, respectively, shall get and set  
 31760 the thread creation stack attributes *stackaddr* and *stacksize* in the *attr* object.

31761 The stack attributes specify the area of storage to be used for the created thread's stack. The base  
 31762 (lowest addressable byte) of the storage shall be *stackaddr*, and the size of the storage shall be  
 31763 *stacksize* bytes. The *stacksize* shall be at least {PTHREAD\_STACK\_MIN}. The *stackaddr* shall be  
 31764 aligned appropriately to be used as a stack; for example, *pthread\_attr\_setstack()* may fail with  
 31765 [EINVAL] if (*stackaddr* & 0x7) is not 0. All pages within the stack described by *stackaddr* and  
 31766 *stacksize* shall be both readable and writable by the thread.

31767 **RETURN VALUE**

31768 Upon successful completion, these functions shall return a value of 0; otherwise, an error  
 31769 number shall be returned to indicate the error.

31770 The *pthread\_attr\_getstack()* function shall store the stack attribute values in *stackaddr* and *stacksize*  
 31771 if successful.

31772 **ERRORS**

31773 The *pthread\_attr\_setstack()* function shall fail if:

31774 [EINVAL] The value of *stacksize* is less than {PTHREAD\_STACK\_MIN} or exceeds an  
 31775 implementation-defined limit.

31776 The *pthread\_attr\_setstack()* function may fail if:

31777 [EINVAL] The value of *stackaddr* does not have proper alignment to be used as a stack, or  
 31778 if (*stackaddr* + *stacksize*) lacks proper alignment.

31779 [EACCES] The stack page(s) described by *stackaddr* and *stacksize* are not both readable  
 31780 and writable by the thread.

31781 These functions shall not return an error code of [EINTR].

31782 **EXAMPLES**

31783 None.

31784 **APPLICATION USAGE**

31785 These functions are appropriate for use by applications in an environment where the stack for a  
 31786 thread must be placed in some particular region of memory.

31787 While it might seem that an application could detect stack overflow by providing a protected  
 31788 page outside the specified stack region, this cannot be done portably. Implementations are free  
 31789 to place the thread's initial stack pointer anywhere within the specified region to accommodate  
 31790 the machine's stack pointer behavior and allocation requirements. Furthermore, on some  
 31791 architectures, such as the IA-64, "overflow" might mean that two separate stack pointers  
 31792 allocated within the region will overlap somewhere in the middle of the region.

31793 **RATIONALE**

31794           None.

31795 **FUTURE DIRECTIONS**

31796           None.

31797 **SEE ALSO**

31798           *pthread\_attr\_init()*, *pthread\_attr\_setdetachstate()*, *pthread\_attr\_setstacksize()*, *pthread\_create()*, the  
31799           Base Definitions volume of IEEE Std 1003.1-200x, <limits.h>, <pthread.h>

31800 **CHANGE HISTORY**

31801           First released in Issue 6. Developed as an XSI extension and brought into the BASE by IEEE  
31802           PASC Interpretation 1003.1 #101.

31803 **NAME**

31804 pthread\_attr\_getstackaddr, pthread\_attr\_setstackaddr — get and set stackaddr attribute

31805 **SYNOPSIS**

31806 THR TSA #include &lt;pthread.h&gt;

```
31807 OB int pthread_attr_getstackaddr(const pthread_attr_t *restrict attr,
31808 void **restrict stackaddr);
31809 int pthread_attr_setstackaddr(pthread_attr_t *attr, void *stackaddr);
31810
```

31811 **DESCRIPTION**

31812 The *pthread\_attr\_getstackaddr()* and *pthread\_attr\_setstackaddr()* functions, respectively, shall get  
 31813 and set the thread creation *stackaddr* attribute in the *attr* object.

31814 The *stackaddr* attribute specifies the location of storage to be used for the created thread's stack.  
 31815 The size of the storage shall be at least {PTHREAD\_STACK\_MIN}.

31816 **RETURN VALUE**

31817 Upon successful completion, *pthread\_attr\_getstackaddr()* and *pthread\_attr\_setstackaddr()* shall  
 31818 return a value of 0; otherwise, an error number shall be returned to indicate the error.

31819 The *pthread\_attr\_getstackaddr()* function stores the *stackaddr* attribute value in *stackaddr* if  
 31820 successful.

31821 **ERRORS**

31822 No errors are defined.

31823 These functions shall not return an error code of [EINTR].

31824 **EXAMPLES**

31825 None.

31826 **APPLICATION USAGE**

31827 The specification of the *stackaddr* attribute presents several ambiguities that make portable use of  
 31828 these interfaces impossible. The description of the single address parameter as a “stack” does  
 31829 not specify a particular relationship between the address and the “stack” implied by that  
 31830 address. For example, the address may be taken as the low memory address of a buffer intended  
 31831 for use as a stack, or it may be taken as the address to be used as the initial stack pointer register  
 31832 value for the new thread. These two are not the same except for a machine on which the stack  
 31833 grows “up” from low memory to high, and on which a “push” operation first stores the value in  
 31834 memory and then increments the stack pointer register. Further, on a machine where the stack  
 31835 grows “down” from high memory to low, interpretation of the address as the “low memory”  
 31836 address requires a determination of the intended size of the stack. IEEE Std 1003.1-200x has  
 31837 introduced the new interfaces *pthread\_attr\_setstack()* and *pthread\_attr\_getstack()* to resolve these  
 31838 ambiguities.

31839 **RATIONALE**

31840 None.

31841 **FUTURE DIRECTIONS**

31842 None.

31843 **SEE ALSO**

31844 *pthread\_attr\_destroy()*, *pthread\_attr\_getdetachstate()*, *pthread\_attr\_getstack()*,  
 31845 *pthread\_attr\_getstacksize()*, *pthread\_attr\_setstack()*, *pthread\_create()*, the Base Definitions volume  
 31846 of IEEE Std 1003.1-200x, <limits.h>, <pthread.h>

## 31847 CHANGE HISTORY

31848 First released in Issue 5. Included for alignment with the POSIX Threads Extension.

## 31849 Issue 6

31850 The *pthread\_attr\_getstackaddr()* and *pthread\_attr\_setstackaddr()* functions are marked as part of  
31851 the Threads and Thread Stack Address Attribute options.

31852 The **restrict** keyword is added to the *pthread\_attr\_getstackaddr()* prototype for alignment with the  
31853 ISO/IEC 9899:1999 standard.

31854 These functions are marked obsolescent.

31855 **NAME**

31856 pthread\_attr\_getstacksize, pthread\_attr\_setstacksize — get and set stacksize attribute

31857 **SYNOPSIS**

31858 THR TSA #include <pthread.h>

```
31859 int pthread_attr_getstacksize(const pthread_attr_t *restrict attr,
31860                             size_t *restrict stacksize);
31861 int pthread_attr_setstacksize(pthread_attr_t *attr, size_t stacksize);
31862
```

31863 **DESCRIPTION**

31864 The *pthread\_attr\_getstacksize()* and *pthread\_attr\_setstacksize()* functions, respectively, shall get  
31865 and set the thread creation *stacksize* attribute in the *attr* object.

31866 The *stacksize* attribute shall define the minimum stack size (in bytes) allocated for the created  
31867 threads stack.

31868 **RETURN VALUE**

31869 Upon successful completion, *pthread\_attr\_getstacksize()* and *pthread\_attr\_setstacksize()* shall  
31870 return a value of 0; otherwise, an error number shall be returned to indicate the error.

31871 The *pthread\_attr\_getstacksize()* function stores the *stacksize* attribute value in *stacksize* if  
31872 successful.

31873 **ERRORS**

31874 The *pthread\_attr\_setstacksize()* function shall fail if:

31875 [EINVAL] The value of *stacksize* is less than {PTHREAD\_STACK\_MIN} or exceeds a  
31876 system-imposed limit.

31877 These functions shall not return an error code of [EINTR].

31878 **EXAMPLES**

31879 None.

31880 **APPLICATION USAGE**

31881 None.

31882 **RATIONALE**

31883 None.

31884 **FUTURE DIRECTIONS**

31885 None.

31886 **SEE ALSO**

31887 *pthread\_attr\_destroy()*, *pthread\_attr\_getstackaddr()*, *pthread\_attr\_getdetachstate()*, *pthread\_create()*,  
31888 the Base Definitions volume of IEEE Std 1003.1-200x, <limits.h>, <pthread.h>

31889 **CHANGE HISTORY**

31890 First released in Issue 5. Included for alignment with the POSIX Threads Extension.

31891 **Issue 6**

31892 The *pthread\_attr\_getstacksize()* and *pthread\_attr\_setstacksize()* functions are marked as part of the  
31893 Threads and Thread Stack Address Attribute options.

31894 The **restrict** keyword is added to the *pthread\_attr\_getstacksize()* prototype for alignment with the  
31895 ISO/IEC 9899:1999 standard.

31896 **NAME**

31897 pthread\_attr\_init — initialize threads attributes object

31898 **SYNOPSIS**

31899 THR #include <pthread.h>

31900 int pthread\_attr\_init(pthread\_attr\_t \*attr);

31901

31902 **DESCRIPTION**

31903 Refer to *pthread\_attr\_destroy()*.



31904 **NAME**

31905 pthread\_attr\_setdetachstate — set detachstate attribute

31906 **SYNOPSIS**

31907 THR #include &lt;pthread.h&gt;

31908 int pthread\_attr\_setdetachstate(pthread\_attr\_t \*attr, int detachstate);

31909

31910 **DESCRIPTION**31911 Refer to *pthread\_attr\_getdetachstate()*.

31912 **NAME**

31913 pthread\_attr\_setguardsize — set thread guardsize attribute

31914 **SYNOPSIS**

31915 XSI #include <pthread.h>

```
31916 int pthread_attr_setguardsize(pthread_attr_t *attr,  
31917 size_t guardsize);
```

31918

31919 **DESCRIPTION**

31920 Refer to *pthread\_attr\_getguardsize()*.

31921 **NAME**31922 pthread\_attr\_setinheritsched — set inheritsched attribute (**REALTIME THREADS**)31923 **SYNOPSIS**

31924 THR TPS #include &lt;pthread.h&gt;

31925 int pthread\_attr\_setinheritsched(pthread\_attr\_t \*attr,  
31926 int inheritsched);

31927

31928 **DESCRIPTION**31929 Refer to *pthread\_attr\_getinheritsched()*.

31930 **NAME**

31931 pthread\_attr\_setschedparam — set schedparam attribute

31932 **SYNOPSIS**

31933 THR #include <pthread.h>

```
31934 int pthread_attr_setschedparam(pthread_attr_t *restrict attr,  
31935                               const struct sched_param *restrict param);  
31936
```

31937 **DESCRIPTION**

31938 Refer to *pthread\_attr\_getschedparam()*.

31939 **NAME**31940 pthread\_attr\_setschedpolicy — set schedpolicy attribute (**REALTIME THREADS**)31941 **SYNOPSIS**

31942 THR TPS #include &lt;pthread.h&gt;

31943 int pthread\_attr\_setschedpolicy(pthread\_attr\_t \*attr, int policy);

31944

31945 **DESCRIPTION**31946 Refer to *pthread\_attr\_getschedpolicy()*.

31947 **NAME**

31948 pthread\_attr\_setscope — set contentionscope attribute (**REALTIME THREADS**)

31949 **SYNOPSIS**

31950 THR TPS #include <pthread.h>

31951 int pthread\_attr\_setscope(pthread\_attr\_t \*attr, int contentionscope);

31952

31953 **DESCRIPTION**

31954 Refer to *pthread\_attr\_getscope()*.

31955 **NAME**

31956 pthread\_attr\_setstack — set stack attribute

31957 **SYNOPSIS**

31958 XSI #include &lt;pthread.h&gt;

31959 int pthread\_attr\_setstack(pthread\_attr\_t \*attr, void \*stackaddr,  
31960 size\_t stacksize);

31961

31962 **DESCRIPTION**31963 Refer to *pthread\_attr\_getstack()*.

31964 **NAME**

31965 pthread\_attr\_setstackaddr — set stackaddr attribute

31966 **SYNOPSIS**

31967 THR TSA #include <pthread.h>

31968 OB int pthread\_attr\_setstackaddr(pthread\_attr\_t \*attr, void \*stackaddr);

31969

31970 **DESCRIPTION**

31971 Refer to *pthread\_attr\_getstackaddr()*.



31972 **NAME**

31973 pthread\_attr\_setstacksize — set stacksize attribute

31974 **SYNOPSIS**

31975 THR TSA #include &lt;pthread.h&gt;

31976 int pthread\_attr\_setstacksize(pthread\_attr\_t \*attr, size\_t stacksize);

31977

31978 **DESCRIPTION**31979 Refer to *pthread\_attr\_getstacksize()*.

## 31980 NAME

31981 pthread\_barrier\_destroy, pthread\_barrier\_init — destroy and initialize a barrier object  
 31982 (ADVANCED REALTIME THREADS)

## 31983 SYNOPSIS

```
31984 THR BAR #include <pthread.h>
31985
31986 int pthread_barrier_destroy(pthread_barrier_t *barrier);
31987 int pthread_barrier_init(pthread_barrier_t *restrict barrier,
31988     const pthread_barrierattr_t *restrict attr, unsigned count);
```

## 31989 DESCRIPTION

31990 The *pthread\_barrier\_destroy()* function shall destroy the barrier referenced by *barrier* and release  
 31991 any resources used by the barrier. The effect of subsequent use of the barrier is undefined until  
 31992 the barrier is reinitialized by another call to *pthread\_barrier\_init()*. An implementation may use  
 31993 this function to set *barrier* to an invalid value. The results are undefined if  
 31994 *pthread\_barrier\_destroy()* is called when any thread is blocked on the barrier, or if this function is  
 31995 called with an uninitialized barrier.

31996 The *pthread\_barrier\_init()* function shall allocate any resources required to use the barrier  
 31997 referenced by *barrier* and shall initialize the barrier with attributes referenced by *attr*. If *attr* is  
 31998 NULL, the default barrier attributes shall be used; the effect is the same as passing the address of  
 31999 a default barrier attributes object. The results are undefined if *pthread\_barrier\_init()* is called  
 32000 when any thread is blocked on the barrier (that is, has not returned from the  
 32001 *pthread\_barrier\_wait()* call). The results are undefined if a barrier is used without first being  
 32002 initialized. The results are undefined if *pthread\_barrier\_init()* is called specifying an already  
 32003 initialized barrier.

32004 The *count* argument specifies the number of threads that must call *pthread\_barrier\_wait()* before  
 32005 any of them successfully return from the call. The value specified by *count* must be greater than  
 32006 zero.

32007 If the *pthread\_barrier\_init()* function fails, the barrier shall not be initialized and the contents of  
 32008 *barrier* are undefined.

32009 Only the object referenced by *barrier* may be used for performing synchronization. The result of  
 32010 referring to copies of that object in calls to *pthread\_barrier\_destroy()* or *pthread\_barrier\_wait()* is  
 32011 undefined.

## 32012 RETURN VALUE

32013 Upon successful completion, these functions shall return zero; otherwise, an error number shall  
 32014 be returned to indicate the error.

## 32015 ERRORS

32016 The *pthread\_barrier\_destroy()* function may fail if:

32017 [EBUSY] The implementation has detected an attempt to destroy a barrier while it is in  
 32018 use (for example, while being used in a *pthread\_barrier\_wait()* call) by another  
 32019 thread.

32020 [EINVAL] The value specified by *barrier* is invalid.

32021 The *pthread\_barrier\_init()* function shall fail if:

32022 [EAGAIN] The system lacks the necessary resources to initialize another barrier.

32023 [EINVAL] The value specified by *count* is equal to zero.

- 32024 [ENOMEM] Insufficient memory exists to initialize the barrier.
- 32025 The *pthread\_barrier\_init()* function may fail if:
- 32026 [EBUSY] The implementation has detected an attempt to reinitialize a barrier while it is  
32027 in use (for example, while being used in a *pthread\_barrier\_wait()* call) by  
32028 another thread.
- 32029 [EINVAL] The value specified by *attr* is invalid.
- 32030 These functions shall not return an error code of [EINTR].
- 32031 **EXAMPLES**
- 32032 None.
- 32033 **APPLICATION USAGE**
- 32034 The *pthread\_barrier\_destroy()* and *pthread\_barrier\_init()* functions are part of the Barriers option  
32035 and need not be provided on all implementations.
- 32036 **RATIONALE**
- 32037 None.
- 32038 **FUTURE DIRECTIONS**
- 32039 None.
- 32040 **SEE ALSO**
- 32041 *pthread\_barrier\_wait()*, the Base Definitions volume of IEEE Std 1003.1-200x, <pthread.h>
- 32042 **CHANGE HISTORY**
- 32043 First released in Issue 6. Derived from IEEE Std 1003.1j-2000.

32044 **NAME**

32045 pthread\_barrier\_init — initialize a barrier object (**ADVANCED REALTIME THREADS**)

32046 **SYNOPSIS**

32047 THR BAR #include <pthread.h>

```
32048 int pthread_barrier_init(pthread_barrier_t *restrict barrier,  
32049                          const pthread_barrierattr_t *restrict attr, unsigned count);
```

32050

32051 **DESCRIPTION**

32052 Refer to *pthread\_barrier\_destroy()*.

32053 **NAME**32054 pthread\_barrier\_wait — synchronize at a barrier (**ADVANCED REALTIME THREADS**)32055 **SYNOPSIS**

32056 THR BAR #include &lt;pthread.h&gt;

32057 int pthread\_barrier\_wait(pthread\_barrier\_t \*barrier);

32058

32059 **DESCRIPTION**

32060 The *pthread\_barrier\_wait()* function shall synchronize participating threads at the barrier  
 32061 referenced by *barrier*. The calling thread shall block until the required number of threads have  
 32062 called *pthread\_barrier\_wait()* specifying the barrier.

32063 When the required number of threads have called *pthread\_barrier\_wait()* specifying the barrier,  
 32064 the constant PTHREAD\_BARRIER\_SERIAL\_THREAD shall be returned to one unspecified  
 32065 thread and zero shall be returned to each of the remaining threads. At this point, the barrier shall  
 32066 be reset to the state it had as a result of the most recent *pthread\_barrier\_init()* function that  
 32067 referenced it.

32068 The constant PTHREAD\_BARRIER\_SERIAL\_THREAD is defined in <pthread.h> and its value  
 32069 shall be distinct from any other value returned by *pthread\_barrier\_wait()*.

32070 The results are undefined if this function is called with an uninitialized barrier.

32071 If a signal is delivered to a thread blocked on a barrier, upon return from the signal handler the  
 32072 thread shall resume waiting at the barrier if the barrier wait has not completed (that is, if the  
 32073 required number of threads have not arrived at the barrier during the execution of the signal  
 32074 handler); otherwise, the thread shall continue as normal from the completed barrier wait. Until  
 32075 the thread in the signal handler returns from it, it is unspecified whether other threads may  
 32076 proceed past the barrier once they have all reached it.

32077 A thread that has blocked on a barrier shall not prevent any unblocked thread that is eligible to  
 32078 use the same processing resources from eventually making forward progress in its execution.  
 32079 Eligibility for processing resources shall be determined by the scheduling policy.

32080 **RETURN VALUE**

32081 Upon successful completion, the *pthread\_barrier\_wait()* function shall return  
 32082 PTHREAD\_BARRIER\_SERIAL\_THREAD for a single (arbitrary) thread synchronized at the  
 32083 barrier and zero for each of the other threads. Otherwise, an error number shall be returned to  
 32084 indicate the error.

32085 **ERRORS**

32086 The *pthread\_barrier\_wait()* function may fail if:

32087 [EINVAL] The value specified by *barrier* does not refer to an initialized barrier object.

32088 This function shall not return an error code of [EINTR].

32089 **EXAMPLES**

32090 None.

32091 **APPLICATION USAGE**

32092 Applications using this function may be subject to priority inversion, as discussed in the Base  
 32093 Definitions volume of IEEE Std 1003.1-200x, Section 3.285, Priority Inversion.

32094 The *pthread\_barrier\_wait()* function is part of the Barriers option and need not be provided on all  
 32095 implementations.

32096 **RATIONALE**

32097       None.

32098 **FUTURE DIRECTIONS**

32099       None.

32100 **SEE ALSO**

32101       *pthread\_barrier\_destroy()*, *pthread\_barrier\_init()*, the Base Definitions volume of  
32102       IEEE Std 1003.1-200x, <pthread.h>

32103 **CHANGE HISTORY**

32104       First released in Issue 6. Derived from IEEE Std 1003.1j-2000.

32105       In the SYNOPSIS, the inclusion of <sys/types.h> is no longer required.

32106 **NAME**

32107 pthread\_barrierattr\_destroy, pthread\_barrierattr\_init — destroy and initialize barrier attributes  
 32108 object (**ADVANCED REALTIME THREADS**)

32109 **SYNOPSIS**

```
32110 THR BAR #include <pthread.h>
```

```
32111 int pthread_barrierattr_destroy(pthread_barrierattr_t *attr);
```

```
32112 int pthread_barrierattr_init(pthread_barrierattr_t *attr);
```

32113

32114 **DESCRIPTION**

32115 The *pthread\_barrierattr\_destroy()* function shall destroy a barrier attributes object. A destroyed  
 32116 *attr* attributes object can be reinitialized using *pthread\_barrierattr\_init()*; the results of otherwise  
 32117 referencing the object after it has been destroyed are undefined. An implementation may cause  
 32118 *pthread\_barrierattr\_destroy()* to set the object referenced by *attr* to an invalid value.

32119 The *pthread\_barrierattr\_init()* function shall initialize a barrier attributes object *attr* with the  
 32120 default value for all of the attributes defined by the implementation.

32121 Results are undefined if *pthread\_barrierattr\_init()* is called specifying an already initialized *attr*  
 32122 attributes object.

32123 After a barrier attributes object has been used to initialize one or more barriers, any function  
 32124 affecting the attributes object (including destruction) shall not affect any previously initialized  
 32125 barrier.

32126 **RETURN VALUE**

32127 If successful, the *pthread\_barrierattr\_destroy()* and *pthread\_barrierattr\_init()* functions shall return  
 32128 zero; otherwise, an error number shall be returned to indicate the error.

32129 **ERRORS**

32130 The *pthread\_barrierattr\_destroy()* function may fail if:

32131 [EINVAL] The value specified by *attr* is invalid.

32132 The *pthread\_barrierattr\_init()* function shall fail if:

32133 [ENOMEM] Insufficient memory exists to initialize the barrier attributes object.

32134 These functions shall not return an error code of [EINTR].

32135 **EXAMPLES**

32136 None.

32137 **APPLICATION USAGE**

32138 The *pthread\_barrierattr\_destroy()* and *pthread\_barrierattr\_init()* functions are part of the Barriers  
 32139 option and need not be provided on all implementations.

32140 **RATIONALE**

32141 None.

32142 **FUTURE DIRECTIONS**

32143 None.

32144 **SEE ALSO**

32145 *pthread\_barrierattr\_getshared()*, *pthread\_barrierattr\_setshared()*, the Base Definitions volume of  
 32146 IEEE Std 1003.1-200x, <pthread.h>.

32147 **CHANGE HISTORY**

32148           First released in Issue 6. Derived from IEEE Std 1003.1j-2000.

32149           In the SYNOPSIS, the inclusion of `<sys/types.h>` is no longer required.



32150 **NAME**

32151 pthread\_barrierattr\_getpshared, pthread\_barrierattr\_setpshared — get and set process-shared  
 32152 attribute of barrier attributes object (**ADVANCED REALTIME THREADS**)

32153 **SYNOPSIS**

32154 THR #include <pthread.h>

```
32155 BAR TSH int pthread_barrierattr_getpshared(
32156           const pthread_barrierattr_t *restrict attr,
32157           int *restrict pshared);
32158 int pthread_barrierattr_setpshared(pthread_barrierattr_t *attr,
32159           int pshared);
32160
```

32161 **DESCRIPTION**

32162 The *pthread\_barrierattr\_getpshared()* function shall obtain the value of the process-shared |  
 32163 attribute from the attributes object referenced by *attr*. The *pthread\_barrierattr\_setpshared()* |  
 32164 function shall set the process-shared attribute in an initialized attributes object referenced by |  
 32165 *attr*. |

32166 The process-shared attribute is set to `PTHREAD_PROCESS_SHARED` to permit a barrier to be |  
 32167 operated upon by any thread that has access to the memory where the barrier is allocated. If the |  
 32168 process-shared attribute is `PTHREAD_PROCESS_PRIVATE`, the barrier shall only be operated |  
 32169 upon by threads created within the same process as the thread that initialized the barrier; if |  
 32170 threads of different processes attempt to operate on such a barrier, the behavior is undefined. |  
 32171 The default value of the attribute shall be `PTHREAD_PROCESS_PRIVATE`. Both constants |  
 32172 `PTHREAD_PROCESS_SHARED` and `PTHREAD_PROCESS_PRIVATE` are defined in |  
 32173 `<pthread.h>`.

32174 Additional attributes, their default values, and the names of the associated functions to get and |  
 32175 set those attribute values are implementation-defined.

32176 **RETURN VALUE**

32177 If successful, the *pthread\_barrierattr\_getpshared()* function shall return zero and store the value of  
 32178 the process-shared attribute of *attr* into the object referenced by the *pshared* parameter.  
 32179 Otherwise, an error number shall be returned to indicate the error.

32180 If successful, the *pthread\_barrierattr\_setpshared()* function shall return zero; otherwise, an error  
 32181 number shall be returned to indicate the error.

32182 **ERRORS**

32183 These functions may fail if:

32184 [EINVAL] The value specified by *attr* is invalid.

32185 The *pthread\_barrierattr\_setpshared()* function may fail if:

32186 [EINVAL] The new value specified for the process-shared attribute is not one of the legal  
 32187 values `PTHREAD_PROCESS_SHARED` or `PTHREAD_PROCESS_PRIVATE`.

32188 These functions shall not return an error code of [EINTR].

32189 **EXAMPLES**

32190           None.

32191 **APPLICATION USAGE**

32192           The *pthread\_barrierattr\_getpshared()* and *pthread\_barrierattr\_setpshared()* functions are part of the  
32193           Barriers option and need not be provided on all implementations.

32194 **RATIONALE**

32195           None.

32196 **FUTURE DIRECTIONS**

32197           None.

32198 **SEE ALSO**

32199           *pthread\_barrier\_init()*, *pthread\_barrierattr\_destroy()*, *pthread\_barrierattr\_init()*, the Base Definitions  
32200           volume of IEEE Std 1003.1-200x, <**pthread.h**>

32201 **CHANGE HISTORY**

32202           First released in Issue 6. Derived from IEEE Std 1003.1j-2000

32203 **NAME**

32204 pthread\_barrierattr\_init — initialize barrier attributes object (**ADVANCED REALTIME**  
32205 **THREADS**)

32206 **SYNOPSIS**

32207 THR BAR #include <pthread.h>

32208 int pthread\_barrierattr\_init(pthread\_barrierattr\_t \*attr);

32209

32210 **DESCRIPTION**

32211 Refer to *pthread\_barrierattr\_destroy()*.

32212 **NAME**

32213 pthread\_barrierattr\_setpshared — set process-shared attribute of barrier attributes object  
32214 (**ADVANCED REALTIME THREADS**)

32215 **SYNOPSIS**

```
32216 THR #include <pthread.h>
```

```
32217 BAR TSH int pthread_barrierattr_setpshared(pthread_barrierattr_t *attr,  
32218 int pshared);
```

32219

32220 **DESCRIPTION**

32221 Refer to *pthread\_barrierattr\_getpshared()*.

32222 **NAME**

32223 pthread\_cancel — cancel execution of a thread

32224 **SYNOPSIS**

32225 THR #include &lt;pthread.h&gt;

32226 int pthread\_cancel(pthread\_t thread);

32227

32228 **DESCRIPTION**

32229 The *pthread\_cancel()* function shall request that *thread* be canceled. The target thread's  
 32230 cancelability state and type determines when the cancelation takes effect. When the cancelation  
 32231 is acted on, the cancelation cleanup handlers for *thread* shall be called. When the last cancelation  
 32232 cleanup handler returns, the thread-specific data destructor functions shall be called for *thread*.  
 32233 When the last destructor function returns, *thread* shall be terminated.

32234 The cancelation processing in the target thread shall run asynchronously with respect to the  
 32235 calling thread returning from *pthread\_cancel()*.

32236 **RETURN VALUE**

32237 If successful, the *pthread\_cancel()* function shall return zero; otherwise, an error number shall be  
 32238 returned to indicate the error.

32239 **ERRORS**32240 The *pthread\_cancel()* function may fail if:

32241 [ESRCH] No thread could be found corresponding to that specified by the given thread  
 32242 ID.

32243 The *pthread\_cancel()* function shall not return an error code of [EINTR].32244 **EXAMPLES**

32245 None.

32246 **APPLICATION USAGE**

32247 None.

32248 **RATIONALE**

32249 Two alternative functions were considered to sending the cancelation notification to a thread.  
 32250 One would be to define a new SIGCANCEL signal that had the cancelation semantics when  
 32251 delivered; the other was to define the new *pthread\_cancel()* function, which would trigger the  
 32252 cancelation semantics.

32253 The advantage of a new signal was that so much of the delivery criteria were identical to that  
 32254 used when trying to deliver a signal that making cancelation notification a signal was seen as  
 32255 consistent. Indeed, many implementations implement cancelation using a special signal. On the  
 32256 other hand, there would be no signal functions that could be used with this signal except  
 32257 *pthread\_kill()*, and the behavior of the delivered cancelation signal would be unlike any  
 32258 previously existing defined signal.

32259 The benefits of a special function include the recognition that this signal would be defined  
 32260 because of the similar delivery criteria and that this is the only common behavior between a  
 32261 cancelation request and a signal. In addition, the cancelation delivery mechanism does not have  
 32262 to be implemented as a signal. There are also strong, if not stronger, parallels with language  
 32263 exception mechanisms than with signals that are potentially obscured if the delivery mechanism  
 32264 is visibly closer to signals.

32265 In the end, it was considered that as there were so many exceptions to the use of the new signal  
 32266 with existing signals functions that it would be misleading. A special function has resolved this

32267 problem. This function was carefully defined so that an implementation wishing to provide the  
32268 cancelation functions on top of signals could do so. The special function also means that  
32269 implementations are not obliged to implement cancelation with signals.

## 32270 FUTURE DIRECTIONS

32271 None.

## 32272 SEE ALSO

32273 *pthread\_exit()*, *pthread\_cond\_wait()*, *pthread\_cond\_timedwait()*, *pthread\_join()*,  
32274 *pthread\_setcancelstate()*, the Base Definitions volume of IEEE Std 1003.1-200x, <pthread.h>

## 32275 CHANGE HISTORY

32276 First released in Issue 5. Included for alignment with the POSIX Threads Extension.

## 32277 Issue 6

32278 The *pthread\_cancel()* function is marked as part of the Threads option.

32279 **NAME**

32280 pthread\_cleanup\_pop, pthread\_cleanup\_push — establish cancelation handlers

32281 **SYNOPSIS**

```
32282 THR #include <pthread.h>
32283 void pthread_cleanup_pop(int execute);
32284 void pthread_cleanup_push(void (*routine)(void*), void *arg);
32285
```

32286 **DESCRIPTION**

32287 The *pthread\_cleanup\_pop()* function shall remove the routine at the top of the calling thread's  
 32288 cancelation cleanup stack and optionally invoke it (if *execute* is non-zero).

32289 The *pthread\_cleanup\_push()* function shall push the specified cancelation cleanup handler *routine*  
 32290 onto the calling thread's cancelation cleanup stack. The cancelation cleanup handler shall be  
 32291 popped from the cancelation cleanup stack and invoked with the argument *arg* when:

- 32292 • The thread exits (that is, calls *pthread\_exit()*).
- 32293 • The thread acts upon a cancelation request.
- 32294 • The thread calls *pthread\_cleanup\_pop()* with a non-zero *execute* argument.

32295 These functions may be implemented as macros. The application shall ensure that they appear  
 32296 as statements, and in pairs within the same lexical scope (that is, the *pthread\_cleanup\_push()*  
 32297 macro may be thought to expand to a token list whose first token is '{' with  
 32298 *pthread\_cleanup\_pop()* expanding to a token list whose last token is the corresponding '}').

32299 The effect of calling *longjmp()* or *siglongjmp()* is undefined if there have been any calls to  
 32300 *pthread\_cleanup\_push()* or *pthread\_cleanup\_pop()* made without the matching call since the jump  
 32301 buffer was filled. The effect of calling *longjmp()* or *siglongjmp()* from inside a cancelation  
 32302 cleanup handler is also undefined unless the jump buffer was also filled in the cancelation  
 32303 cleanup handler.

32304 **RETURN VALUE**32305 The *pthread\_cleanup\_push()* and *pthread\_cleanup\_pop()* functions shall not return a value.32306 **ERRORS**

32307 No errors are defined.

32308 These functions shall not return an error code of [EINTR].

32309 **EXAMPLES**

32310 The following is an example using thread primitives to implement a cancelable, writers-priority  
 32311 read-write lock:

```
32312 typedef struct {
32313     pthread_mutex_t lock;
32314     pthread_cond_t rcond,
32315     wcond;
32316     int lock_count; /* < 0 .. Held by writer. */
32317                   /* > 0 .. Held by lock_count readers. */
32318                   /* = 0 .. Held by nobody. */
32319     int waiting_writers; /* Count of waiting writers. */
32320 } rwlock;
32321 void
32322 waiting_reader_cleanup(void *arg)
32323 {
```

```
32324         rwlock *l;
32325         l = (rwlock *) arg;
32326         pthread_mutex_unlock(&l->lock);
32327     }
32328 void
32329 lock_for_read(rwlock *l)
32330 {
32331     pthread_mutex_lock(&l->lock);
32332     pthread_cleanup_push(waiting_reader_cleanup, l);
32333     while ((l->lock_count < 0) && (l->waiting_writers != 0))
32334         pthread_cond_wait(&l->rcond, &l->lock);
32335     l->lock_count++;
32336     /*
32337      * Note the pthread_cleanup_pop executes
32338      * waiting_reader_cleanup.
32339      */
32340     pthread_cleanup_pop(1);
32341 }
32342 void
32343 release_read_lock(rwlock *l)
32344 {
32345     pthread_mutex_lock(&l->lock);
32346     if (--l->lock_count == 0)
32347         pthread_cond_signal(&l->wcond);
32348     pthread_mutex_unlock(l);
32349 }
32350 void
32351 waiting_writer_cleanup(void *arg)
32352 {
32353     rwlock *l;
32354     l = (rwlock *) arg;
32355     if ((--l->waiting_writers == 0) && (l->lock_count >= 0)) {
32356         /*
32357          * This only happens if we have been canceled.
32358          */
32359         pthread_cond_broadcast(&l->wcond);
32360     }
32361     pthread_mutex_unlock(&l->lock);
32362 }
32363 void
32364 lock_for_write(rwlock *l)
32365 {
32366     pthread_mutex_lock(&l->lock);
32367     l->waiting_writers++;
32368     pthread_cleanup_push(waiting_writer_cleanup, l);
32369     while (l->lock_count != 0)
32370         pthread_cond_wait(&l->wcond, &l->lock);
32371     l->lock_count = -1;
32372     /*
```



```
32373         * Note the pthread_cleanup_pop executes
32374         * waiting_writer_cleanup.
32375         */
32376         pthread_cleanup_pop(1);
32377     }

32378     void
32379     release_write_lock(rwlock *l)
32380     {
32381         pthread_mutex_lock(&l->lock);
32382         l->lock_count = 0;
32383         if (l->waiting_writers == 0)
32384             pthread_cond_broadcast(&l->rcond)
32385         else
32386             pthread_cond_signal(&l->wcond);
32387         pthread_mutex_unlock(&l->lock);
32388     }

32389     /*
32390     * This function is called to initialize the read/write lock.
32391     */
32392     void
32393     initialize_rwlock(rwlock *l)
32394     {
32395         pthread_mutex_init(&l->lock, pthread_mutexattr_default);
32396         pthread_cond_init(&l->wcond, pthread_condattr_default);
32397         pthread_cond_init(&l->rcond, pthread_condattr_default);
32398         l->lock_count = 0;
32399         l->waiting_writers = 0;
32400     }

32401     reader_thread()
32402     {
32403         lock_for_read(&lock);
32404         pthread_cleanup_push(release_read_lock, &lock);
32405         /*
32406         * Thread has read lock.
32407         */
32408         pthread_cleanup_pop(1);
32409     }

32410     writer_thread()
32411     {
32412         lock_for_write(&lock);
32413         pthread_cleanup_push(release_write_lock, &lock);
32414         /*
32415         * Thread has write lock.
32416         */
32417         pthread_cleanup_pop(1);
32418     }
```

32419 **APPLICATION USAGE**

32420 The two routines that push and pop cancelation cleanup handlers, *pthread\_cleanup\_push()* and  
32421 *pthread\_cleanup\_pop()*, can be thought of as left and right parentheses. They always need to be  
32422 matched.

32423 **RATIONALE**

32424 The restriction that the two routines that push and pop cancelation cleanup handlers,  
32425 *pthread\_cleanup\_push()* and *pthread\_cleanup\_pop()*, have to appear in the same lexical scope  
32426 allows for efficient macro or compiler implementations and efficient storage management. A  
32427 sample implementation of these routines as macros might look like this:

```
32428 #define pthread_cleanup_push(rtn,arg) { \  
32429     struct _pthread_handler_rec __cleanup_handler, **__head; \  
32430     __cleanup_handler.rtn = rtn; \  
32431     __cleanup_handler.arg = arg; \  
32432     (void) pthread_getspecific(_pthread_handler_key, &__head); \  
32433     __cleanup_handler.next = *__head; \  
32434     *__head = &__cleanup_handler;  
  
32435 #define pthread_cleanup_pop(ex) \  
32436     *__head = __cleanup_handler.next; \  
32437     if (ex) (*__cleanup_handler.rtn)(__cleanup_handler.arg); \  
32438 }
```

32439 A more ambitious implementation of these routines might do even better by allowing the  
32440 compiler to note that the cancelation cleanup handler is a constant and can be expanded inline.

32441 This volume of IEEE Std 1003.1-200x currently leaves unspecified the effect of calling *longjmp()*  
32442 from a signal handler executing in a POSIX System Interfaces function. If an implementation  
32443 wants to allow this and give the programmer reasonable behavior, the *longjmp()* function has to  
32444 call all cancelation cleanup handlers that have been pushed but not popped since the time  
32445 *setjmp()* was called.

32446 Consider a multi-threaded function called by a thread that uses signals. If a signal were  
32447 delivered to a signal handler during the operation of *qsort()* and that handler were to call  
32448 *longjmp()* (which, in turn, did *not* call the cancelation cleanup handlers) the helper threads  
32449 created by the *qsort()* function would not be canceled. Instead, they would continue to execute  
32450 and write into the argument array even though the array might have been popped off of the  
32451 stack.

32452 Note that the specified cleanup handling mechanism is especially tied to the C language and,  
32453 while the requirement for a uniform mechanism for expressing cleanup is language-  
32454 independent, the mechanism used in other languages may be quite different. In addition, this  
32455 mechanism is really only necessary due to the lack of a real exception mechanism in the C  
32456 language, which would be the ideal solution.

32457 There is no notion of a cancelation cleanup-safe function. If an application has no cancelation  
32458 points in its signal handlers, blocks any signal whose handler may have cancelation points while  
32459 calling async-unsafe functions, or disables cancelation while calling async-unsafe functions, all  
32460 functions may be safely called from cancelation cleanup routines.

32461 **FUTURE DIRECTIONS**

32462 None.

32463 **SEE ALSO**

32464 *pthread\_cancel()*, *pthread\_setcancelstate()*, the Base Definitions volume of IEEE Std 1003.1-200x,  
32465 <pthread.h>

32466 **CHANGE HISTORY**

32467 First released in Issue 5. Included for alignment with the POSIX Threads Extension.

32468 **Issue 6**

32469 The *pthread\_cleanup\_pop()* and *pthread\_cleanup\_push()* functions are marked as part of the  
32470 Threads option.

32471 The APPLICATION USAGE section is added.

32472 The DESCRIPTION is updated to avoid use of the term “must” for application requirements.

## 32473 NAME

32474 pthread\_cond\_broadcast, pthread\_cond\_signal — broadcast or signal a condition

## 32475 SYNOPSIS

```
32476 THR #include <pthread.h>
```

```
32477 int pthread_cond_broadcast(pthread_cond_t *cond);
```

```
32478 int pthread_cond_signal(pthread_cond_t *cond);
```

32479

## 32480 DESCRIPTION

32481 These functions shall unblock threads blocked on a condition variable. |

32482 The *pthread\_cond\_broadcast()* function shall unblock all threads currently blocked on the |  
32483 specified condition variable *cond*.

32484 The *pthread\_cond\_signal()* function shall unblock at least one of the threads that are blocked on |  
32485 the specified condition variable *cond* (if any threads are blocked on *cond*).

32486 If more than one thread is blocked on a condition variable, the scheduling policy shall determine |  
32487 the order in which threads are unblocked. When each thread unblocked as a result of a |  
32488 *pthread\_cond\_broadcast()* or *pthread\_cond\_signal()* returns from its call to *pthread\_cond\_wait()* or |  
32489 *pthread\_cond\_timedwait()*, the thread shall own the mutex with which it called |  
32490 *pthread\_cond\_wait()* or *pthread\_cond\_timedwait()*. The thread(s) that are unblocked shall contend |  
32491 for the mutex according to the scheduling policy (if applicable), and as if each had called |  
32492 *pthread\_mutex\_lock()*.

32493 The *pthread\_cond\_broadcast()* or *pthread\_cond\_signal()* functions may be called by a thread |  
32494 whether or not it currently owns the mutex that threads calling *pthread\_cond\_wait()* or |  
32495 *pthread\_cond\_timedwait()* have associated with the condition variable during their waits; |  
32496 however, if predictable scheduling behavior is required, then that mutex shall be locked by the |  
32497 thread calling *pthread\_cond\_broadcast()* or *pthread\_cond\_signal()*.

32498 The *pthread\_cond\_broadcast()* and *pthread\_cond\_signal()* functions shall have no effect if there are |  
32499 no threads currently blocked on *cond*. |

## 32500 RETURN VALUE

32501 If successful, the *pthread\_cond\_broadcast()* and *pthread\_cond\_signal()* functions shall return zero; |  
32502 otherwise, an error number shall be returned to indicate the error.

## 32503 ERRORS

32504 The *pthread\_cond\_broadcast()* and *pthread\_cond\_signal()* function may fail if:

32505 [EINVAL] The value *cond* does not refer to an initialized condition variable.

32506 These functions shall not return an error code of [EINTR].

## 32507 EXAMPLES

32508 None.

## 32509 APPLICATION USAGE

32510 The *pthread\_cond\_broadcast()* function is used whenever the shared-variable state has been |  
32511 changed in a way that more than one thread can proceed with its task. Consider a single |  
32512 producer/multiple consumer problem, where the producer can insert multiple items on a list |  
32513 that is accessed one item at a time by the consumers. By calling the *pthread\_cond\_broadcast()* |  
32514 function, the producer would notify all consumers that might be waiting, and thereby the |  
32515 application would receive more throughput on a multi-processor. In addition, |  
32516 *pthread\_cond\_broadcast()* makes it easier to implement a read-write lock. The |  
32517 *pthread\_cond\_broadcast()* function is needed in order to wake up all waiting readers when a

32518 writer releases its lock. Finally, the two-phase commit algorithm can use this broadcast function  
32519 to notify all clients of an impending transaction commit.

32520 It is not safe to use the *pthread\_cond\_signal()* function in a signal handler that is invoked  
32521 asynchronously. Even if it were safe, there would still be a race between the test of the Boolean  
32522 *pthread\_cond\_wait()* that could not be efficiently eliminated.

32523 Mutexes and condition variables are thus not suitable for releasing a waiting thread by signaling  
32524 from code running in a signal handler.

### 32525 RATIONALE

#### 32526 Multiple Awakenings by Condition Signal

32527 On a multi-processor, it may be impossible for an implementation of *pthread\_cond\_signal()* to  
32528 avoid the unblocking of more than one thread blocked on a condition variable. For example,  
32529 consider the following partial implementation of *pthread\_cond\_wait()* and *pthread\_cond\_signal()*,  
32530 executed by two threads in the order given. One thread is trying to wait on the condition  
32531 variable, another is concurrently executing *pthread\_cond\_signal()*, while a third thread is already  
32532 waiting.

```
32533 pthread_cond_wait(mutex, cond):
32534     value = cond->value; /* 1 */
32535     pthread_mutex_unlock(mutex); /* 2 */
32536     pthread_mutex_lock(cond->mutex); /* 10 */
32537     if (value == cond->value) { /* 11 */
32538         me->next_cond = cond->waiter;
32539         cond->waiter = me;
32540         pthread_mutex_unlock(cond->mutex);
32541         unable_to_run(me);
32542     } else
32543         pthread_mutex_unlock(cond->mutex); /* 12 */
32544     pthread_mutex_lock(mutex); /* 13 */

32545 pthread_cond_signal(cond):
32546     pthread_mutex_lock(cond->mutex); /* 3 */
32547     cond->value++; /* 4 */
32548     if (cond->waiter) { /* 5 */
32549         sleeper = cond->waiter; /* 6 */
32550         cond->waiter = sleeper->next_cond; /* 7 */
32551         able_to_run(sleeper); /* 8 */
32552     }
32553     pthread_mutex_unlock(cond->mutex); /* 9 */
```

32554 The effect is that more than one thread can return from its call to *pthread\_cond\_wait()* or  
32555 *pthread\_cond\_timedwait()* as a result of one call to *pthread\_cond\_signal()*. This effect is called  
32556 “spurious wakeup”. Note that the situation is self-correcting in that the number of threads that  
32557 are so awakened is finite; for example, the next thread to call *pthread\_cond\_wait()* after the  
32558 sequence of events above blocks.

32559 While this problem could be resolved, the loss of efficiency for a fringe condition that occurs  
32560 only rarely is unacceptable, especially given that one has to check the predicate associated with a  
32561 condition variable anyway. Correcting this problem would unnecessarily reduce the degree of  
32562 concurrency in this basic building block for all higher-level synchronization operations.

32563 An added benefit of allowing spurious wakeups is that applications are forced to code a  
32564 predicate-testing-loop around the condition wait. This also makes the application tolerate

32565           superfluous condition broadcasts or signals on the same condition variable that may be coded in  
32566           some other part of the application. The resulting applications are thus more robust. Therefore,  
32567           IEEE Std 1003.1-200x explicitly documents that spurious wakeups may occur.

## 32568 **FUTURE DIRECTIONS**

32569           None.

## 32570 **SEE ALSO**

32571           *pthread\_cond\_destroy()*, *pthread\_cond\_timedwait()*, *pthread\_cond\_wait()*, the Base Definitions  
32572           volume of IEEE Std 1003.1-200x, <**pthread.h**>

## 32573 **CHANGE HISTORY**

32574           First released in Issue 5. Included for alignment with the POSIX Threads Extension.

## 32575 **Issue 6**

32576           The *pthread\_cond\_broadcast()* and *pthread\_cond\_signal()* functions are marked as part of the  
32577           Threads option.

32578           The APPLICATION USAGE section is added.

32579 **NAME**

32580 pthread\_cond\_destroy, pthread\_cond\_init — destroy and initialize condition variables

32581 **SYNOPSIS**

32582 THR #include &lt;pthread.h&gt;

```

32583 int pthread_cond_destroy(pthread_cond_t *cond);
32584 int pthread_cond_init(pthread_cond_t *restrict cond,
32585     const pthread_condattr_t *restrict attr);
32586 pthread_cond_t cond = PTHREAD_COND_INITIALIZER;
32587

```

32588 **DESCRIPTION**

32589 The *pthread\_cond\_destroy()* function shall destroy the given condition variable specified by *cond*;  
 32590 the object becomes, in effect, uninitialized. An implementation may cause *pthread\_cond\_destroy()*  
 32591 to set the object referenced by *cond* to an invalid value. A destroyed condition variable object can  
 32592 be reinitialized using *pthread\_cond\_init()*; the results of otherwise referencing the object after it  
 32593 has been destroyed are undefined.

32594 It shall be safe to destroy an initialized condition variable upon which no threads are currently  
 32595 blocked. Attempting to destroy a condition variable upon which other threads are currently  
 32596 blocked results in undefined behavior.

32597 The *pthread\_cond\_init()* function shall initialize the condition variable referenced by *cond* with  
 32598 attributes referenced by *attr*. If *attr* is NULL, the default condition variable attributes shall be  
 32599 used; the effect is the same as passing the address of a default condition variable attributes  
 32600 object. Upon successful initialization, the state of the condition variable shall become initialized.

32601 Only *cond* itself may be used for performing synchronization. The result of referring to copies of  
 32602 *cond* in calls to *pthread\_cond\_wait()*, *pthread\_cond\_timedwait()*, *pthread\_cond\_signal()*,  
 32603 *pthread\_cond\_broadcast()*, and *pthread\_cond\_destroy()* is undefined.

32604 Attempting to initialize an already initialized condition variable results in undefined behavior.

32605 In cases where default condition variable attributes are appropriate, the macro  
 32606 PTHREAD\_COND\_INITIALIZER can be used to initialize condition variables that are statically  
 32607 allocated. The effect shall be equivalent to dynamic initialization by a call to *pthread\_cond\_init()*  
 32608 with parameter *attr* specified as NULL, except that no error checks are performed.

32609 **RETURN VALUE**

32610 If successful, the *pthread\_cond\_destroy()* and *pthread\_cond\_init()* functions shall return zero;  
 32611 otherwise, an error number shall be returned to indicate the error.

32612 The [EBUSY] and [EINVAL] error checks, if implemented, shall act as if they were performed  
 32613 immediately at the beginning of processing for the function and caused an error return prior to  
 32614 modifying the state of the condition variable specified by *cond*.

32615 **ERRORS**

32616 The *pthread\_cond\_destroy()* function may fail if:

32617 [EBUSY] The implementation has detected an attempt to destroy the object referenced  
 32618 by *cond* while it is referenced (for example, while being used in a  
 32619 *pthread\_cond\_wait()* or *pthread\_cond\_timedwait()*) by another thread.

32620 [EINVAL] The value specified by *cond* is invalid.

32621 The *pthread\_cond\_init()* function shall fail if:

32622 [EAGAIN] The system lacked the necessary resources (other than memory) to initialize  
 32623 another condition variable.

32624 [ENOMEM] Insufficient memory exists to initialize the condition variable.

32625 The `pthread_cond_init()` function may fail if:

32626 [EBUSY] The implementation has detected an attempt to reinitialize the object |  
32627 referenced by `cond`, a previously initialized, but not yet destroyed, condition  
32628 variable.

32629 [EINVAL] The value specified by `attr` is invalid.

32630 These functions shall not return an error code of [EINTR].

### 32631 EXAMPLES

32632 A condition variable can be destroyed immediately after all the threads that are blocked on it are  
32633 awakened. For example, consider the following code:

```
32634 struct list {
32635     pthread_mutex_t lm;
32636     ...
32637 }
32638 struct elt {
32639     key k;
32640     int busy;
32641     pthread_cond_t notbusy;
32642     ...
32643 }
32644 /* Find a list element and reserve it. */
32645 struct elt *
32646 list_find(struct list *lp, key k)
32647 {
32648     struct elt *ep;
32649     pthread_mutex_lock(&lp->lm);
32650     while ((ep = find_elt(l, k) != NULL) && ep->busy)
32651         pthread_cond_wait(&ep->notbusy, &lp->lm);
32652     if (ep != NULL)
32653         ep->busy = 1;
32654     pthread_mutex_unlock(&lp->lm);
32655     return(ep);
32656 }
32657 delete_elt(struct list *lp, struct elt *ep)
32658 {
32659     pthread_mutex_lock(&lp->lm);
32660     assert(ep->busy);
32661     ... remove ep from list ...
32662     ep->busy = 0; /* Paranoid. */
32663     (A) pthread_cond_broadcast(&ep->notbusy);
32664     pthread_mutex_unlock(&lp->lm);
32665     (B) pthread_cond_destroy(&rp->notbusy);
32666     free(ep);
32667 }
```

32668 In this example, the condition variable and its list element may be freed (line B) immediately  
32669 after all threads waiting for it are awakened (line A), since the mutex and the code ensure that no  
32670 other thread can touch the element to be deleted.



32671 **APPLICATION USAGE**

32672 None.

32673 **RATIONALE**32674 See *pthread\_mutex\_init()*; a similar rationale applies to condition variables.32675 **FUTURE DIRECTIONS**

32676 None.

32677 **SEE ALSO**32678 *pthread\_cond\_broadcast()*, *pthread\_cond\_signal()*, *pthread\_cond\_timedwait()*, *pthread\_cond\_wait()*,  
32679 the Base Definitions volume of IEEE Std 1003.1-200x, <pthread.h>32680 **CHANGE HISTORY**

32681 First released in Issue 5. Included for alignment with the POSIX Threads Extension.

32682 **Issue 6**32683 The *pthread\_cond\_destroy()* and *pthread\_cond\_init()* functions are marked as part of the Threads  
32684 option.

32685 IEEE PASC Interpretation 1003.1c #34 is applied, updating the DESCRIPTION.

32686 The **restrict** keyword is added to the *pthread\_cond\_init()* prototype for alignment with the  
32687 ISO/IEC 9899:1999 standard.

32688 **NAME**

32689 pthread\_cond\_init — initialize condition variables

32690 **SYNOPSIS**

32691 THR #include <pthread.h>

32692 int pthread\_cond\_init(pthread\_cond\_t \*restrict cond,

32693 const pthread\_condattr\_t \*restrict attr);

32694 pthread\_cond\_t cond = PTHREAD\_COND\_INITIALIZER;

32695

32696 **DESCRIPTION**

32697 Refer to *pthread\_cond\_destroy()*.

32698 **NAME**

32699 pthread\_cond\_signal — signal a condition

32700 **SYNOPSIS**

32701 THR #include &lt;pthread.h&gt;

32702 int pthread\_cond\_signal(pthread\_cond\_t \*cond);

32703

32704 **DESCRIPTION**32705 Refer to *pthread\_cond\_broadcast()*.

## 32706 NAME

32707 pthread\_cond\_timedwait, pthread\_cond\_wait — wait on a condition

## 32708 SYNOPSIS

```
32709 THR #include <pthread.h>
32710 int pthread_cond_timedwait(pthread_cond_t *restrict cond,
32711 pthread_mutex_t *restrict mutex,
32712 const struct timespec *restrict abstime);
32713 int pthread_cond_wait(pthread_cond_t *restrict cond,
32714 pthread_mutex_t *restrict mutex);
32715
```

## 32716 DESCRIPTION

32717 The *pthread\_cond\_timedwait()* and *pthread\_cond\_wait()* functions shall block on a condition |  
 32718 variable. They shall be called with *mutex* locked by the calling thread or undefined behavior |  
 32719 results.

32720 These functions atomically release *mutex* and cause the calling thread to block on the condition  
 32721 variable *cond*; atomically here means “atomically with respect to access by another thread to the  
 32722 mutex and then the condition variable”. That is, if another thread is able to acquire the mutex  
 32723 after the about-to-block thread has released it, then a subsequent call to *pthread\_cond\_broadcast()*  
 32724 or *pthread\_cond\_signal()* in that thread shall behave as if it were issued after the about-to-block  
 32725 thread has blocked.

32726 Upon successful return, the mutex shall have been locked and shall be owned by the calling |  
 32727 thread. |

32728 When using condition variables there is always a Boolean predicate involving shared variables  
 32729 associated with each condition wait that is true if the thread should proceed. Spurious wakeups  
 32730 from the *pthread\_cond\_timedwait()* or *pthread\_cond\_wait()* functions may occur. Since the return  
 32731 from *pthread\_cond\_timedwait()* or *pthread\_cond\_wait()* does not imply anything about the value  
 32732 of this predicate, the predicate should be re-evaluated upon such return.

32733 The effect of using more than one mutex for concurrent *pthread\_cond\_timedwait()* or  
 32734 *pthread\_cond\_wait()* operations on the same condition variable is undefined; that is, a condition  
 32735 variable becomes bound to a unique mutex when a thread waits on the condition variable, and  
 32736 this (dynamic) binding shall end when the wait returns.

32737 A condition wait (whether timed or not) is a cancelation point. When the cancelability enable  
 32738 state of a thread is set to PTHREAD\_CANCEL\_DEFERRED, a side effect of acting upon a  
 32739 cancelation request while in a condition wait is that the mutex is (in effect) re-acquired before  
 32740 calling the first cancelation cleanup handler. The effect is as if the thread were unblocked,  
 32741 allowed to execute up to the point of returning from the call to *pthread\_cond\_timedwait()* or  
 32742 *pthread\_cond\_wait()*, but at that point notices the cancelation request and instead of returning to  
 32743 the caller of *pthread\_cond\_timedwait()* or *pthread\_cond\_wait()*, starts the thread cancelation  
 32744 activities, which includes calling cancelation cleanup handlers.

32745 A thread that has been unblocked because it has been canceled while blocked in a call to  
 32746 *pthread\_cond\_timedwait()* or *pthread\_cond\_wait()* shall not consume any condition signal that  
 32747 may be directed concurrently at the condition variable if there are other threads blocked on the  
 32748 condition variable.

32749 The *pthread\_cond\_timedwait()* function shall be equivalent to *pthread\_cond\_wait()*, except that an |  
 32750 error is returned if the absolute time specified by *abstime* passes (that is, system time equals or  
 32751 exceeds *abstime*) before the condition *cond* is signaled or broadcasted, or if the absolute time

32752 CS

32753 specified by *abstime* has already been passed at the time of the call. If the Clock Selection option  
 32754 is supported, the condition variable shall have a clock attribute which specifies the clock that  
 32755 shall be used to measure the time specified by the *abstime* argument. When such timeouts occur,  
 32756 *pthread\_cond\_timedwait()* shall nonetheless release and re-acquire the mutex referenced by *mutex*.  
 32757 The *pthread\_cond\_timedwait()* function is also a cancellation point.

32758 If a signal is delivered to a thread waiting for a condition variable, upon return from the signal  
 32759 handler the thread resumes waiting for the condition variable as if it was not interrupted, or it  
 32760 shall return zero due to spurious wakeup.

#### 32761 RETURN VALUE

32762 Except in the case of [ETIMEDOUT], all these error checks shall act as if they were performed  
 32763 immediately at the beginning of processing for the function and shall cause an error return, in  
 32764 effect, prior to modifying the state of the mutex specified by *mutex* or the condition variable  
 32765 specified by *cond*.

32766 Upon successful completion, a value of zero shall be returned; otherwise, an error number shall  
 32767 be returned to indicate the error.

#### 32768 ERRORS

32769 The *pthread\_cond\_timedwait()* function shall fail if:

32770 [ETIMEDOUT] The time specified by *abstime* to *pthread\_cond\_timedwait()* has passed.

32771 The *pthread\_cond\_timedwait()* and *pthread\_cond\_wait()* functions may fail if:

32772 [EINVAL] The value specified by *cond*, *mutex*, or *abstime* is invalid.

32773 [EINVAL] Different mutexes were supplied for concurrent *pthread\_cond\_timedwait()* or  
 32774 *pthread\_cond\_wait()* operations on the same condition variable.

32775 [EPERM] The mutex was not owned by the current thread at the time of the call.

32776 These functions shall not return an error code of [EINTR].

#### 32777 EXAMPLES

32778 None.

#### 32779 APPLICATION USAGE

32780 None.

#### 32781 RATIONALE

##### 32782 Condition Wait Semantics

32783 It is important to note that when *pthread\_cond\_wait()* and *pthread\_cond\_timedwait()* return  
 32784 without error, the associated predicate may still be false. Similarly, when  
 32785 *pthread\_cond\_timedwait()* returns with the timeout error, the associated predicate may be true  
 32786 due to an unavoidable race between the expiration of the timeout and the predicate state change.

32787 Some implementations, particularly on a multi-processor, may sometimes cause multiple  
 32788 threads to wake up when the condition variable is signaled simultaneously on different  
 32789 processors.

32790 In general, whenever a condition wait returns, the thread has to re-evaluate the predicate  
 32791 associated with the condition wait to determine whether it can safely proceed, should wait  
 32792 again, or should declare a timeout. A return from the wait does not imply that the associated  
 32793 predicate is either true or false.

32794 It is thus recommended that a condition wait be enclosed in the equivalent of a “while loop”  
 32795 that checks the predicate.

32796 **Timed Wait Semantics**

32797 An absolute time measure was chosen for specifying the timeout parameter for two reasons.  
32798 First, a relative time measure can be easily implemented on top of a function that specifies  
32799 absolute time, but there is a race condition associated with specifying an absolute timeout on top  
32800 of a function that specifies relative timeouts. For example, assume that `clock_gettime()` returns  
32801 the current time and `cond_relative_timed_wait()` uses relative timeouts:

```
32802 clock_gettime(CLOCK_REALTIME, &now)  
32803 reltime = sleep_til_this_absolute_time -now;  
32804 cond_relative_timed_wait(c, m, &reltime);
```

32805 If the thread is preempted between the first statement and the last statement, the thread blocks  
32806 for too long. Blocking, however, is irrelevant if an absolute timeout is used. An absolute timeout  
32807 also need not be recomputed if it is used multiple times in a loop, such as that enclosing a  
32808 condition wait.

32809 For cases when the system clock is advanced discontinuously by an operator, it is expected that  
32810 implementations process any timed wait expiring at an intervening time as if that time had  
32811 actually occurred.

32812 **Cancelation and Condition Wait**

32813 A condition wait, whether timed or not, is a cancelation point. That is, the functions  
32814 `pthread_cond_wait()` or `pthread_cond_timedwait()` are points where a pending (or concurrent)  
32815 cancelation request is noticed. The reason for this is that an indefinite wait is possible at these  
32816 points—whatever event is being waited for, even if the program is totally correct, might never  
32817 occur; for example, some input data being awaited might never be sent. By making condition  
32818 wait a cancelation point, the thread can be canceled and perform its cancelation cleanup handler  
32819 even though it may be stuck in some indefinite wait.

32820 A side effect of acting on a cancelation request while a thread is blocked on a condition variable  
32821 is to re-acquire the mutex before calling any of the cancelation cleanup handlers. This is done in  
32822 order to ensure that the cancelation cleanup handler is executed in the same state as the critical  
32823 code that lies both before and after the call to the condition wait function. This rule is also  
32824 required when interfacing to POSIX threads from languages, such as Ada or C++, which may  
32825 choose to map cancelation onto a language exception; this rule ensures that each exception  
32826 handler guarding a critical section can always safely depend upon the fact that the associated  
32827 mutex has already been locked regardless of exactly where within the critical section the  
32828 exception was raised. Without this rule, there would not be a uniform rule that exception  
32829 handlers could follow regarding the lock, and so coding would become very cumbersome.

32830 Therefore, since *some* statement has to be made regarding the state of the lock when a  
32831 cancelation is delivered during a wait, a definition has been chosen that makes application  
32832 coding most convenient and error free.

32833 When acting on a cancelation request while a thread is blocked on a condition variable, the  
32834 implementation is required to ensure that the thread does not consume any condition signals  
32835 directed at that condition variable if there are any other threads waiting on that condition  
32836 variable. This rule is specified in order to avoid deadlock conditions that could occur if these two  
32837 independent requests (one acting on a thread and the other acting on the condition variable)  
32838 were not processed independently.

32839 **Performance of Mutexes and Condition Variables**

32840 Mutexes are expected to be locked only for a few instructions. This practice is almost  
 32841 automatically enforced by the desire of programmers to avoid long serial regions of execution  
 32842 (which would reduce total effective parallelism).

32843 When using mutexes and condition variables, one tries to ensure that the usual case is to lock the  
 32844 mutex, access shared data, and unlock the mutex. Waiting on a condition variable should be a  
 32845 relatively rare situation. For example, when implementing a read-write lock, code that acquires a  
 32846 read-lock typically needs only to increment the count of readers (under mutual-exclusion) and  
 32847 return. The calling thread would actually wait on the condition variable only when there is  
 32848 already an active writer. So the efficiency of a synchronization operation is bounded by the cost  
 32849 of mutex lock/unlock and not by condition wait. Note that in the usual case there is no context  
 32850 switch.

32851 This is not to say that the efficiency of condition waiting is unimportant. Since there needs to be  
 32852 at least one context switch per Ada rendezvous, the efficiency of waiting on a condition variable  
 32853 is important. The cost of waiting on a condition variable should be little more than the minimal  
 32854 cost for a context switch plus the time to unlock and lock the mutex.

32855 **Features of Mutexes and Condition Variables**

32856 It had been suggested that the mutex acquisition and release be decoupled from condition wait.  
 32857 This was rejected because it is the combined nature of the operation that, in fact, facilitates  
 32858 realtime implementations. Those implementations can atomically move a high-priority thread  
 32859 between the condition variable and the mutex in a manner that is transparent to the caller. This  
 32860 can prevent extra context switches and provide more deterministic acquisition of a mutex when  
 32861 the waiting thread is signaled. Thus, fairness and priority issues can be dealt with directly by the  
 32862 scheduling discipline. Furthermore, the current condition wait operation matches existing  
 32863 practice.

32864 **Scheduling Behavior of Mutexes and Condition Variables**

32865 Synchronization primitives that attempt to interfere with scheduling policy by specifying an  
 32866 ordering rule are considered undesirable. Threads waiting on mutexes and condition variables  
 32867 are selected to proceed in an order dependent upon the scheduling policy rather than in some  
 32868 fixed order (for example, FIFO or priority). Thus, the scheduling policy determines which  
 32869 thread(s) are awakened and allowed to proceed.

32870 **Timed Condition Wait**

32871 The *pthread\_cond\_timedwait()* function allows an application to give up waiting for a particular  
 32872 condition after a given amount of time. An example of its use follows:

```
32873 (void) pthread_mutex_lock(&t.mn);
32874     t.waiters++;
32875     clock_gettime(CLOCK_REALTIME, &ts);
32876     ts.tv_sec += 5;
32877     rc = 0;
32878     while (! mypredicate(&t) && rc == 0)
32879         rc = pthread_cond_timedwait(&t.cond, &t.mn, &ts);
32880     t.waiters--;
32881     if (rc == 0) setmystate(&t);
32882 (void) pthread_mutex_unlock(&t.mn);
```

32883 By making the timeout parameter absolute, it does not need to be recomputed each time the  
32884 program checks its blocking predicate. If the timeout was relative, it would have to be  
32885 recomputed before each call. This would be especially difficult since such code would need to  
32886 take into account the possibility of extra wakeups that result from extra broadcasts or signals on  
32887 the condition variable that occur before either the predicate is true or the timeout is due.

32888 **FUTURE DIRECTIONS**

32889 None.

32890 **SEE ALSO**

32891 *pthread\_cond\_signal()*, *pthread\_cond\_broadcast()*, the Base Definitions volume of  
32892 IEEE Std 1003.1-200x, <pthread.h>

32893 **CHANGE HISTORY**

32894 First released in Issue 5. Included for alignment with the POSIX Threads Extension.

32895 **Issue 6**

32896 The *pthread\_cond\_timedwait()* and *pthread\_cond\_wait()* functions are marked as part of the  
32897 Threads option.

32898 The Open Group Corrigendum U021/9 is applied, correcting the prototype for the  
32899 *pthread\_cond\_wait()* function.

32900 The DESCRIPTION is updated for alignment with IEEE Std 1003.1j-2000 by adding semantics for  
32901 the Clock Selection option.

32902 The ERRORS section has an additional case for [EPERM] in response to IEEE PASC  
32903 Interpretation 1003.1c #28.

32904 The **restrict** keyword is added to the *pthread\_cond\_timedwait()* and *pthread\_cond\_wait()*  
32905 prototypes for alignment with the ISO/IEC 9899:1999 standard.



32906 **NAME**

32907 pthread\_cond\_wait — wait on a condition

32908 **SYNOPSIS**

32909 THR #include &lt;pthread.h&gt;

32910 int pthread\_cond\_wait(pthread\_cond\_t \*restrict cond,  
32911 pthread\_mutex\_t \*restrict mutex);

32912

32913 **DESCRIPTION**32914 Refer to *pthread\_cond\_timedwait()*.

## 32915 NAME

32916 pthread\_condattr\_destroy, pthread\_condattr\_init — destroy and initialize condition variable  
32917 attributes object

## 32918 SYNOPSIS

```
32919 THR #include <pthread.h>
```

```
32920 int pthread_condattr_destroy(pthread_condattr_t *attr);
```

```
32921 int pthread_condattr_init(pthread_condattr_t *attr);
```

32922

## 32923 DESCRIPTION

32924 The *pthread\_condattr\_destroy()* function shall destroy a condition variable attributes object; the  
32925 object becomes, in effect, uninitialized. An implementation may cause *pthread\_condattr\_destroy()*  
32926 to set the object referenced by *attr* to an invalid value. A destroyed *attr* attributes object can be  
32927 reinitialized using *pthread\_condattr\_init()*; the results of otherwise referencing the object after it  
32928 has been destroyed are undefined.

32929 The *pthread\_condattr\_init()* function shall initialize a condition variable attributes object *attr* with  
32930 the default value for all of the attributes defined by the implementation.

32931 Results are undefined if *pthread\_condattr\_init()* is called specifying an already initialized *attr*  
32932 attributes object.

32933 After a condition variable attributes object has been used to initialize one or more condition  
32934 variables, any function affecting the attributes object (including destruction) shall not affect any  
32935 previously initialized condition variables.

32936 This volume of IEEE Std 1003.1-200x requires two attributes, the *clock* attribute and the *process-*  
32937 *shared* attribute.

32938 Additional attributes, their default values, and the names of the associated functions to get and  
32939 set those attribute values are implementation-defined.

## 32940 RETURN VALUE

32941 If successful, the *pthread\_condattr\_destroy()* and *pthread\_condattr\_init()* functions shall return  
32942 zero; otherwise, an error number shall be returned to indicate the error.

## 32943 ERRORS

32944 The *pthread\_condattr\_destroy()* function may fail if:

32945 [EINVAL] The value specified by *attr* is invalid.

32946 The *pthread\_condattr\_init()* function shall fail if:

32947 [ENOMEM] Insufficient memory exists to initialize the condition variable attributes object.

32948 These functions shall not return an error code of [EINTR].

## 32949 EXAMPLES

32950 None.

## 32951 APPLICATION USAGE

32952 None.

## 32953 RATIONALE

32954 See *pthread\_attr\_init()* and *pthread\_mutex\_init()*.

32955 A *process-shared* attribute has been defined for condition variables for the same reason it has been  
32956 defined for mutexes.

32957 **FUTURE DIRECTIONS**

32958           None.

32959 **SEE ALSO**32960           *pthread\_cond\_destroy()*, *pthread\_condattr\_getpshared()*, *pthread\_create()*, *pthread\_mutex\_destroy()*,  
32961           the Base Definitions volume of IEEE Std 1003.1-200x, <**pthread.h**>32962 **CHANGE HISTORY**

32963           First released in Issue 5. Included for alignment with the POSIX Threads Extension.

32964 **Issue 6**32965           The *pthread\_condattr\_destroy()* and *pthread\_condattr\_init()* functions are marked as part of the  
32966           Threads option.

32967 **NAME**

32968 pthread\_condattr\_getclock, pthread\_condattr\_setclock — get and set the clock selection  
 32969 condition variable attribute (**ADVANCED REALTIME**)

32970 **SYNOPSIS**

32971 THR CS #include <pthread.h>

```
32972 int pthread_condattr_getclock(const pthread_condattr_t *restrict attr,
32973 clockid_t *restrict clock_id);
32974 int pthread_condattr_setclock(pthread_condattr_t *attr,
32975 clockid_t clock_id);
32976
```

32977 **DESCRIPTION**

32978 The *pthread\_condattr\_getclock()* function shall obtain the value of the *clock* attribute from the  
 32979 attributes object referenced by *attr*. The *pthread\_condattr\_setclock()* function shall set the *clock*  
 32980 attribute in an initialized attributes object referenced by *attr*. If *pthread\_condattr\_setclock()* is  
 32981 called with a *clock\_id* argument that refers to a CPU-time clock, the call shall fail.

32982 The *clock* attribute is the clock ID of the clock that shall be used to measure the timeout service of  
 32983 *pthread\_cond\_timedwait()*. The default value of the *clock* attribute shall refer to the system clock.

32984 **RETURN VALUE**

32985 If successful, the *pthread\_condattr\_getclock()* function shall return zero and store the value of the  
 32986 clock attribute of *attr* into the object referenced by the *clock\_id* argument. Otherwise, an error  
 32987 number shall be returned to indicate the error.

32988 If successful, the *pthread\_condattr\_setclock()* function shall return zero; otherwise, an error  
 32989 number shall be returned to indicate the error.

32990 **ERRORS**

32991 These functions may fail if:

32992 [EINVAL] The value specified by *attr* is invalid.

32993 The *pthread\_condattr\_setclock()* function may fail if:

32994 [EINVAL] The value specified by *clock\_id* does not refer to a known clock, or is a CPU-  
 32995 time clock.

32996 These functions shall not return an error code of [EINTR].

32997 **EXAMPLES**

32998 None.

32999 **APPLICATION USAGE**

33000 None.

33001 **RATIONALE**

33002 None.

33003 **FUTURE DIRECTIONS**

33004 None.

33005 **SEE ALSO**

33006 *pthread\_cond\_init()*, *pthread\_cond\_timedwait()*, *pthread\_condattr\_destroy()*,  
 33007 *pthread\_condattr\_getshared()* (on page 1536), *pthread\_condattr\_init()*,  
 33008 *pthread\_condattr\_setshared()* (on page 1540), *pthread\_create()*, *pthread\_mutex\_init()*, the Base  
 33009 Definitions volume of IEEE Std 1003.1-200x, <pthread.h>

33010 **CHANGE HISTORY**

33011 First released in Issue 6. Derived from IEEE Std 1003.1j-2000.

33012 **NAME**

33013 pthread\_condattr\_getpshared, pthread\_condattr\_setpshared — get and set the process-shared  
 33014 condition variable attributes

33015 **SYNOPSIS**

```
33016 THR TSH #include <pthread.h>
33017
33017 int pthread_condattr_getpshared(const pthread_condattr_t *restrict attr,
33018 int *restrict pshared);
33019 int pthread_condattr_setpshared(pthread_condattr_t *attr,
33020 int pshared);
33021
```

33022 **DESCRIPTION**

33023 The *pthread\_condattr\_getpshared()* function shall obtain the value of the *process-shared* attribute  
 33024 from the attributes object referenced by *attr*. The *pthread\_condattr\_setpshared()* function shall set  
 33025 the *process-shared* attribute in an initialized attributes object referenced by *attr*.

33026 The *process-shared* attribute is set to `PTHREAD_PROCESS_SHARED` to permit a condition  
 33027 variable to be operated upon by any thread that has access to the memory where the condition  
 33028 variable is allocated, even if the condition variable is allocated in memory that is shared by  
 33029 multiple processes. If the *process-shared* attribute is `PTHREAD_PROCESS_PRIVATE`, the  
 33030 condition variable shall only be operated upon by threads created within the same process as the  
 33031 thread that initialized the condition variable; if threads of differing processes attempt to operate  
 33032 on such a condition variable, the behavior is undefined. The default value of the attribute is  
 33033 `PTHREAD_PROCESS_PRIVATE`.

33034 **RETURN VALUE**

33035 If successful, the *pthread\_condattr\_setpshared()* function shall return zero; otherwise, an error  
 33036 number shall be returned to indicate the error.

33037 If successful, the *pthread\_condattr\_getpshared()* function shall return zero and store the value of  
 33038 the *process-shared* attribute of *attr* into the object referenced by the *pshared* parameter. Otherwise,  
 33039 an error number shall be returned to indicate the error.

33040 **ERRORS**

33041 The *pthread\_condattr\_getpshared()* and *pthread\_condattr\_setpshared()* functions may fail if:

33042 [EINVAL] The value specified by *attr* is invalid.

33043 The *pthread\_condattr\_setpshared()* function may fail if:

33044 [EINVAL] The new value specified for the attribute is outside the range of legal values  
 33045 for that attribute.

33046 These functions shall not return an error code of [EINTR].

33047 **EXAMPLES**

33048 None.

33049 **APPLICATION USAGE**

33050 None.

33051 **RATIONALE**

33052 None.

33053 **FUTURE DIRECTIONS**

33054 None.

33055 **SEE ALSO**

33056 *pthread\_create()*, *pthread\_cond\_destroy()*, *pthread\_condattr\_destroy()*, *pthread\_mutex\_destroy()*, the  
33057 Base Definitions volume of IEEE Std 1003.1-200x, <pthread.h>

33058 **CHANGE HISTORY**

33059 First released in Issue 5. Included for alignment with the POSIX Threads Extension.

33060 **Issue 6**

33061 The *pthread\_condattr\_getpshared()* and *pthread\_condattr\_setpshared()* functions are marked as part  
33062 of the Threads and Thread Process-Shared Synchronization options.

33063 The **restrict** keyword is added to the *pthread\_condattr\_getpshared()* prototype for alignment with  
33064 the ISO/IEC 9899:1999 standard.

33065 **NAME**

33066 pthread\_condattr\_init — initialize condition variable attributes object

33067 **SYNOPSIS**

33068 THR #include <pthread.h>

33069 int pthread\_condattr\_init(pthread\_condattr\_t \*attr);

33070

33071 **DESCRIPTION**

33072 Refer to *pthread\_condattr\_destroy()*.



33073 **NAME**

33074 pthread\_condattr\_setclock — set the clock selection condition variable attribute

33075 **SYNOPSIS**

33076 THR CS #include &lt;pthread.h&gt;

33077 int pthread\_condattr\_setclock(pthread\_condattr\_t \*attr,  
33078 clockid\_t clock\_id);

33079

33080 **DESCRIPTION**33081 Refer to *pthread\_condattr\_getclock()*.

33082 **NAME**

33083 pthread\_condattr\_setpshared — set the process-shared condition variable attributes

33084 **SYNOPSIS**

33085 THR TSH #include <pthread.h>

```
33086 int pthread_condattr_setpshared(pthread_condattr_t *attr,  
33087 int pshared);
```

33088

33089 **DESCRIPTION**

33090 Refer to *pthread\_condattr\_getpshared()*.

## 33091 NAME

33092 pthread\_create — thread creation

## 33093 SYNOPSIS

33094 THR #include &lt;pthread.h&gt;

```
33095 int pthread_create(pthread_t *restrict thread,
33096                  const pthread_attr_t *restrict attr,
33097                  void *(*start_routine)(void*), void *restrict arg);
33098
```

## 33099 DESCRIPTION

33100 The *pthread\_create()* function shall create a new thread, with attributes specified by *attr*, within a |  
 33101 process. If *attr* is NULL, the default attributes shall be used. If the attributes specified by *attr* are |  
 33102 modified later, the thread's attributes shall not be affected. Upon successful completion, |  
 33103 *pthread\_create()* shall store the ID of the created thread in the location referenced by *thread*.

33104 The thread is created executing *start\_routine* with *arg* as its sole argument. If the *start\_routine* |  
 33105 returns, the effect shall be as if there was an implicit call to *pthread\_exit()* using the return value |  
 33106 of *start\_routine* as the exit status. Note that the thread in which *main()* was originally invoked |  
 33107 differs from this. When it returns from *main()*, the effect shall be as if there was an implicit call |  
 33108 to *exit()* using the return value of *main()* as the exit status.

33109 The signal state of the new thread shall be initialized as follows:

- 33110 • The signal mask shall be inherited from the creating thread.
- 33111 • The set of signals pending for the new thread shall be empty.

33112 The floating-point environment shall be inherited from the creating thread. |

33113 If *pthread\_create()* fails, no new thread is created and the contents of the location referenced by |  
 33114 *thread* are undefined.

33115 TCT If `_POSIX_THREAD_CPUTIME` is defined, the new thread shall have a CPU-time clock |  
 33116 accessible, and the initial value of this clock shall be set to zero.

## 33117 RETURN VALUE

33118 If successful, the *pthread\_create()* function shall return zero; otherwise, an error number shall be |  
 33119 returned to indicate the error.

## 33120 ERRORS

33121 The *pthread\_create()* function shall fail if:

33122 [EAGAIN] The system lacked the necessary resources to create another thread, or the |  
 33123 system-imposed limit on the total number of threads in a process |  
 33124 PTHREAD\_THREADS\_MAX would be exceeded.

33125 [EINVAL] The value specified by *attr* is invalid.

33126 [EPERM] The caller does not have appropriate permission to set the required |  
 33127 scheduling parameters or scheduling policy.

33128 The *pthread\_create()* function shall not return an error code of [EINTR].

33129 **EXAMPLES**

33130 None.

33131 **APPLICATION USAGE**

33132 None.

33133 **RATIONALE**

33134 A suggested alternative to *pthread\_create()* would be to define two separate operations: create  
 33135 and start. Some applications would find such behavior more natural. Ada, in particular,  
 33136 separates the “creation” of a task from its “activation”.

33137 Splitting the operation was rejected by the standard developers for many reasons:

- 33138 • The number of calls required to start a thread would increase from one to two and thus place  
 33139 an additional burden on applications that do not require the additional synchronization. The  
 33140 second call, however, could be avoided by the additional complication of a start-up state  
 33141 attribute.
- 33142 • An extra state would be introduced: “created but not started”. This would require the  
 33143 standard to specify the behavior of the thread operations when the target has not yet started  
 33144 executing.
- 33145 • For those applications that require such behavior, it is possible to simulate the two separate  
 33146 steps with the facilities that are currently provided. The *start\_routine()* can synchronize by  
 33147 waiting on a condition variable that is signaled by the start operation.

33148 An Ada implementor can choose to create the thread at either of two points in the Ada program:  
 33149 when the task object is created, or when the task is activated (generally at a “begin”). If the first  
 33150 approach is adopted, the *start\_routine()* needs to wait on a condition variable to receive the  
 33151 order to begin “activation”. The second approach requires no such condition variable or extra  
 33152 synchronization. In either approach, a separate Ada task control block would need to be created  
 33153 when the task object is created to hold rendezvous queues, and so on.

33154 An extension of the preceding model would be to allow the state of the thread to be modified  
 33155 between the create and start. This would allow the thread attributes object to be eliminated. This  
 33156 has been rejected because:

- 33157 • All state in the thread attributes object has to be able to be set for the thread. This would  
 33158 require the definition of functions to modify thread attributes. There would be no reduction  
 33159 in the number of function calls required to set up the thread. In fact, for an application that  
 33160 creates all threads using identical attributes, the number of function calls required to set up  
 33161 the threads would be dramatically increased. Use of a thread attributes object permits the  
 33162 application to make one set of attribute setting function calls. Otherwise, the set of attribute  
 33163 setting function calls needs to be made for each thread creation.
- 33164 • Depending on the implementation architecture, functions to set thread state would require  
 33165 kernel calls, or for other implementation reasons would not be able to be implemented as  
 33166 macros, thereby increasing the cost of thread creation.
- 33167 • The ability for applications to segregate threads by class would be lost.

33168 Another suggested alternative uses a model similar to that for process creation, such as “thread  
 33169 fork”. The fork semantics would provide more flexibility and the “create” function can be  
 33170 implemented simply by doing a thread fork followed immediately by a call to the desired “start  
 33171 routine” for the thread. This alternative has these problems:

- 33172 • For many implementations, the entire stack of the calling thread would need to be  
 33173 duplicated, since in many architectures there is no way to determine the size of the calling  
 33174 frame.

33175           • Efficiency is reduced since at least some part of the stack has to be copied, even though in  
33176           most cases the thread never needs the copied context, since it merely calls the desired start  
33177           routine.

33178 **FUTURE DIRECTIONS**

33179           None.

33180 **SEE ALSO**

33181           *fork()*, *pthread\_exit()*, *pthread\_join()*, the Base Definitions volume of IEEE Std 1003.1-200x,  
33182           <pthread.h>

33183 **CHANGE HISTORY**

33184           First released in Issue 5. Included for alignment with the POSIX Threads Extension.

33185 **Issue 6**

33186           The *pthread\_create()* function is marked as part of the Threads option.

33187           The following new requirements on POSIX implementations derive from alignment with the  
33188           Single UNIX Specification:

33189           • The [EPERM] mandatory error condition is added.

33190           The thread CPU-time clock semantics are added for alignment with IEEE Std 1003.1d-1999.

33191           The **restrict** keyword is added to the *pthread\_create()* prototype for alignment with the  
33192           ISO/IEC 9899:1999 standard. |

33193           The DESCRIPTION is updated to make it explicit that the floating-point environment is |  
33194           inherited from the creating thread. |

33195 **NAME**

33196 pthread\_detach — detach a thread

33197 **SYNOPSIS**

33198 THR #include &lt;pthread.h&gt;

33199 int pthread\_detach(pthread\_t thread);

33200

33201 **DESCRIPTION**

33202 The *pthread\_detach()* function shall indicate to the implementation that storage for the thread  
33203 *thread* can be reclaimed when that thread terminates. If *thread* has not terminated,  
33204 *pthread\_detach()* shall not cause it to terminate. The effect of multiple *pthread\_detach()* calls on  
33205 the same target thread is unspecified.

33206 **RETURN VALUE**

33207 If the call succeeds, *pthread\_detach()* shall return 0; otherwise, an error number shall be returned  
33208 to indicate the error.

33209 **ERRORS**33210 The *pthread\_detach()* function shall fail if:

33211 [EINVAL] The implementation has detected that the value specified by *thread* does not  
33212 refer to a joinable thread.

33213 [ESRCH] No thread could be found corresponding to that specified by the given thread  
33214 ID.

33215 The *pthread\_detach()* function shall not return an error code of [EINTR].33216 **EXAMPLES**

33217 None.

33218 **APPLICATION USAGE**

33219 None.

33220 **RATIONALE**

33221 The *pthread\_join()* or *pthread\_detach()* functions should eventually be called for every thread that  
33222 is created so that storage associated with the thread may be reclaimed.

33223 It has been suggested that a “detach” function is not necessary; the *detachstate* thread creation  
33224 attribute is sufficient, since a thread need never be dynamically detached. However, need arises  
33225 in at least two cases:

33226 1. In a cancellation handler for a *pthread\_join()* it is nearly essential to have a *pthread\_detach()*  
33227 function in order to detach the thread on which *pthread\_join()* was waiting. Without it, it  
33228 would be necessary to have the handler do another *pthread\_join()* to attempt to detach the  
33229 thread, which would both delay the cancellation processing for an unbounded period and  
33230 introduce a new call to *pthread\_join()*, which might itself need a cancellation handler. A  
33231 dynamic detach is nearly essential in this case.

33232 2. In order to detach the “initial thread” (as may be desirable in processes that set up server  
33233 threads).

33234 **FUTURE DIRECTIONS**

33235 None.

33236 **SEE ALSO**

33237 *pthread\_join()*, the Base Definitions volume of IEEE Std 1003.1-200x, <pthread.h>

33238 **CHANGE HISTORY**

33239 First released in Issue 5. Included for alignment with the POSIX Threads Extension.

33240 **Issue 6**

33241 The *pthread\_detach()* function is marked as part of the Threads option.

33242 **NAME**

33243 pthread\_equal — compare thread IDs

33244 **SYNOPSIS**

33245 THR #include &lt;pthread.h&gt;

33246 int pthread\_equal(pthread\_t t1, pthread\_t t2);

33247

33248 **DESCRIPTION**33249 This function shall compare the thread IDs *t1* and *t2*.33250 **RETURN VALUE**33251 The *pthread\_equal()* function shall return a non-zero value if *t1* and *t2* are equal; otherwise, zero  
33252 shall be returned.33253 If either *t1* or *t2* are not valid thread IDs, the behavior is undefined.33254 **ERRORS**

33255 No errors are defined.

33256 The *pthread\_equal()* function shall not return an error code of [EINTR].33257 **EXAMPLES**

33258 None.

33259 **APPLICATION USAGE**

33260 None.

33261 **RATIONALE**33262 Implementations may choose to define a thread ID as a structure. This allows additional  
33263 flexibility and robustness over using an **int**. For example, a thread ID could include a sequence  
33264 number that allows detection of “dangling IDs” (copies of a thread ID that has been detached).  
33265 Since the C language does not support comparison on structure types, the *pthread\_equal()*  
33266 function is provided to compare thread IDs.33267 **FUTURE DIRECTIONS**

33268 None.

33269 **SEE ALSO**33270 *pthread\_create()*, *pthread\_self()*, the Base Definitions volume of IEEE Std 1003.1-200x, <pthread.h>33271 **CHANGE HISTORY**

33272 First released in Issue 5. Included for alignment with the POSIX Threads Extension.

33273 **Issue 6**33274 The *pthread\_equal()* function is marked as part of the Threads option.



33275 **NAME**

33276 pthread\_exit — thread termination

33277 **SYNOPSIS**

33278 THR #include &lt;pthread.h&gt;

33279 void pthread\_exit(void \*value\_ptr);

33280

33281 **DESCRIPTION**

33282 The *pthread\_exit()* function shall terminate the calling thread and make the value *value\_ptr*  
 33283 available to any successful join with the terminating thread. Any cancellation cleanup handlers  
 33284 that have been pushed and not yet popped shall be popped in the reverse order that they were  
 33285 pushed and then executed. After all cancellation cleanup handlers have been executed, if the  
 33286 thread has any thread-specific data, appropriate destructor functions shall be called in an  
 33287 unspecified order. Thread termination does not release any application visible process resources,  
 33288 including, but not limited to, mutexes and file descriptors, nor does it perform any process-level  
 33289 cleanup actions, including, but not limited to, calling any *atexit()* routines that may exist.

33290 An implicit call to *pthread\_exit()* is made when a thread other than the thread in which *main()*  
 33291 was first invoked returns from the start routine that was used to create it. The function's return  
 33292 value shall serve as the thread's exit status.

33293 The behavior of *pthread\_exit()* is undefined if called from a cancellation cleanup handler or  
 33294 destructor function that was invoked as a result of either an implicit or explicit call to  
 33295 *pthread\_exit()*.

33296 After a thread has terminated, the result of access to local (auto) variables of the thread is  
 33297 undefined. Thus, references to local variables of the exiting thread should not be used for the  
 33298 *pthread\_exit()* *value\_ptr* parameter value.

33299 The process shall exit with an exit status of 0 after the last thread has been terminated. The  
 33300 behavior shall be as if the implementation called *exit()* with a zero argument at thread  
 33301 termination time.

33302 **RETURN VALUE**33303 The *pthread\_exit()* function cannot return to its caller.33304 **ERRORS**

33305 No errors are defined.

33306 **EXAMPLES**

33307 None.

33308 **APPLICATION USAGE**

33309 None.

33310 **RATIONALE**

33311 The normal mechanism by which a thread terminates is to return from the routine that was  
 33312 specified in the *pthread\_create()* call that started it. The *pthread\_exit()* function provides the  
 33313 capability for a thread to terminate without requiring a return from the start routine of that  
 33314 thread, thereby providing a function analogous to *exit()*.

33315 Regardless of the method of thread termination, any cancellation cleanup handlers that have  
 33316 been pushed and not yet popped are executed, and the destructors for any existing thread-  
 33317 specific data are executed. This volume of IEEE Std 1003.1-200x requires that cancellation  
 33318 cleanup handlers be popped and called in order. After all cancellation cleanup handlers have  
 33319 been executed, thread-specific data destructors are called, in an unspecified order, for each item  
 33320 of thread-specific data that exists in the thread. This ordering is necessary because cancellation

- 33321 cleanup handlers may rely on thread-specific data.
- 33322 As the meaning of the status is determined by the application (except when the thread has been
- 33323 canceled, in which case it is PTHREAD\_CANCELED), the implementation has no idea what an
- 33324 illegal status value is, which is why no address error checking is done.
- 33325 **FUTURE DIRECTIONS**
- 33326 None.
- 33327 **SEE ALSO**
- 33328 *exit()*, *pthread\_create()*, *pthread\_join()*, the Base Definitions volume of IEEE Std 1003.1-200x,
- 33329 **<pthread.h>**
- 33330 **CHANGE HISTORY**
- 33331 First released in Issue 5. Included for alignment with the POSIX Threads Extension.
- 33332 **Issue 6**
- 33333 The *pthread\_exit()* function is marked as part of the Threads option.

33334 **NAME**

33335 pthread\_getconcurrency, pthread\_setconcurrency — get and set level of concurrency

33336 **SYNOPSIS**

33337 XSI #include &lt;pthread.h&gt;

33338 int pthread\_getconcurrency(void);

33339 int pthread\_setconcurrency(int new\_level);

33340

33341 **DESCRIPTION**

33342 Unbound threads in a process may or may not be required to be simultaneously active. By  
 33343 default, the threads implementation ensures that a sufficient number of threads are active so that  
 33344 the process can continue to make progress. While this conserves system resources, it may not  
 33345 produce the most effective level of concurrency.

33346 The *pthread\_setconcurrency()* function allows an application to inform the threads  
 33347 implementation of its desired concurrency level, *new\_level*. The actual level of concurrency  
 33348 provided by the implementation as a result of this function call is unspecified.

33349 If *new\_level* is zero, it causes the implementation to maintain the concurrency level at its  
 33350 discretion as if *pthread\_setconcurrency()* had never been called.

33351 The *pthread\_getconcurrency()* function shall return the value set by a previous call to the  
 33352 *pthread\_setconcurrency()* function. If the *pthread\_setconcurrency()* function was not previously  
 33353 called, this function shall return zero to indicate that the implementation is maintaining the  
 33354 concurrency level.

33355 A call to *pthread\_setconcurrency()* shall inform the implementation of its desired concurrency |  
 33356 level. The implementation shall use this as a hint, not a requirement. |

33357 If an implementation does not support multiplexing of user threads on top of several kernel- |  
 33358 scheduled entities, the *pthread\_setconcurrency()* and *pthread\_getconcurrency()* functions are |  
 33359 provided for source code compatibility but they shall have no effect when called. To maintain |  
 33360 the function semantics, the *new\_level* parameter is saved when *pthread\_setconcurrency()* is called  
 33361 so that a subsequent call to *pthread\_getconcurrency()* shall return the same value.

33362 **RETURN VALUE**

33363 If successful, the *pthread\_setconcurrency()* function shall return zero; otherwise, an error number  
 33364 shall be returned to indicate the error.

33365 The *pthread\_getconcurrency()* function shall always return the concurrency level set by a previous  
 33366 call to *pthread\_setconcurrency()*. If the *pthread\_setconcurrency()* function has never been called,  
 33367 *pthread\_getconcurrency()* shall return zero.

33368 **ERRORS**33369 The *pthread\_setconcurrency()* function shall fail if:33370 [EINVAL] The value specified by *new\_level* is negative.33371 [EAGAIN] The value specific by *new\_level* would cause a system resource to be exceeded.

33372 These functions shall not return an error code of [EINTR].

33373 **EXAMPLES**

33374           None.

33375 **APPLICATION USAGE**

33376           Use of these functions changes the state of the underlying concurrency upon which the  
33377           application depends. Library developers are advised to not use the *pthread\_getconcurrency()* and  
33378           *pthread\_setconcurrency()* functions since their use may conflict with an applications use of these  
33379           functions.

33380 **RATIONALE**

33381           None.

33382 **FUTURE DIRECTIONS**

33383           None.

33384 **SEE ALSO**

33385           The Base Definitions volume of IEEE Std 1003.1-200x, <pthread.h>

33386 **CHANGE HISTORY**

33387           First released in Issue 5.

33388 **NAME**

33389 pthread\_getcpuclockid — access a thread CPU-time clock (**ADVANCED REALTIME**  
33390 **THREADS**)

33391 **SYNOPSIS**

```
33392 THR TCT #include <pthread.h>
```

```
33393 #include <time.h>
```

```
33394 int pthread_getcpuclockid(pthread_t thread_id, clockid_t *clock_id);
```

33395

33396 **DESCRIPTION**

33397 The *pthread\_getcpuclockid()* function shall return in *clock\_id* the clock ID of the CPU-time clock of  
33398 the thread specified by *thread\_id*, if the thread specified by *thread\_id* exists.

33399 **RETURN VALUE**

33400 Upon successful completion, *pthread\_getcpuclockid()* shall return zero; otherwise, an error  
33401 number shall be returned to indicate the error.

33402 **ERRORS**

33403 The *pthread\_getcpuclockid()* function may fail if:

33404 [ESRCH] The value specified by *thread\_id* does not refer to an existing thread.

33405 **EXAMPLES**

33406 None.

33407 **APPLICATION USAGE**

33408 The *pthread\_getcpuclockid()* function is part of the Thread CPU-Time Clocks option and need not  
33409 be provided on all implementations.

33410 **RATIONALE**

33411 None.

33412 **FUTURE DIRECTIONS**

33413 None.

33414 **SEE ALSO**

33415 *clock\_getcpuclockid()*, *clock\_getres()*, *timer\_create()*, the Base Definitions volume of  
33416 IEEE Std 1003.1-200x, <pthread.h>, <time.h>

33417 **CHANGE HISTORY**

33418 First released in Issue 6. Derived from IEEE Std 1003.1d-1999.

33419 In the SYNOPSIS, the inclusion of <sys/types.h> is no longer required.

## 33420 NAME

33421 pthread\_getschedparam, pthread\_setschedparam — dynamic thread scheduling parameters  
 33422 access (**REALTIME THREADS**)

## 33423 SYNOPSIS

33424 THR TPS #include <pthread.h>

```
33425 int pthread_getschedparam(pthread_t thread, int *restrict policy,
33426 struct sched_param *restrict param);
```

```
33427 int pthread_setschedparam(pthread_t thread, int policy,
33428 const struct sched_param *param);
```

33429

## 33430 DESCRIPTION

33431 The *pthread\_getschedparam()* and *pthread\_setschedparam()* functions shall, respectively, get and set |  
 33432 the scheduling policy and parameters of individual threads within a multi-threaded process to |  
 33433 be retrieved and set. For SCHED\_FIFO and SCHED\_RR, the only required member of the |  
 33434 **sched\_param** structure is the priority *sched\_priority*. For SCHED\_OTHER, the affected |  
 33435 scheduling parameters are implementation-defined.

33436 The *pthread\_getschedparam()* function shall retrieve the scheduling policy and scheduling |  
 33437 parameters for the thread whose thread ID is given by *thread* and shall store those values in |  
 33438 *policy* and *param*, respectively. The priority value returned from *pthread\_getschedparam()* shall be |  
 33439 the value specified by the most recent *pthread\_setschedparam()*, *pthread\_setschedprio()*, or |  
 33440 *pthread\_create()* call affecting the target thread. It shall not reflect any temporary adjustments to |  
 33441 its priority as a result of any priority inheritance or ceiling functions. The *pthread\_setschedparam()* |  
 33442 function shall set the scheduling policy and associated scheduling parameters for the thread |  
 33443 whose thread ID is given by *thread* to the policy and associated parameters provided in *policy* |  
 33444 and *param*, respectively.

33445 The *policy* parameter may have the value SCHED\_OTHER, SCHED\_FIFO, or SCHED\_RR. The |  
 33446 scheduling parameters for the SCHED\_OTHER policy are implementation-defined. The |  
 33447 SCHED\_FIFO and SCHED\_RR policies shall have a single scheduling parameter, *priority*.

33448 TSP If \_POSIX\_THREAD\_SPORADIC\_SERVER is defined, then the *policy* argument may have the |  
 33449 value SCHED\_SPORADIC, with the exception for the *pthread\_setschedparam()* function that if the |  
 33450 scheduling policy was not SCHED\_SPORADIC at the time of the call, it is implementation- |  
 33451 defined whether the function is supported; in other words, the implementation need not allow |  
 33452 the application to dynamically change the scheduling policy to SCHED\_SPORADIC. The |  
 33453 sporadic server scheduling policy has the associated parameters *sched\_ss\_low\_priority*, |  
 33454 *sched\_ss\_repl\_period*, *sched\_ss\_init\_budget*, *sched\_priority*, and *sched\_ss\_max\_repl*. The specified |  
 33455 *sched\_ss\_repl\_period* shall be greater than or equal to the specified *sched\_ss\_init\_budget* for the |  
 33456 function to succeed; if it is not, then the function shall fail. The value of *sched\_ss\_max\_repl* shall |  
 33457 be within the inclusive range [1,{SS\_REPL\_MAX}] for the function to succeed; if not, the function |  
 33458 shall fail.

33459 If the *pthread\_setschedparam()* function fails, the scheduling parameters shall not be changed for |  
 33460 the target thread. |

## 33461 RETURN VALUE

33462 If successful, the *pthread\_getschedparam()* and *pthread\_setschedparam()* functions shall return zero; |  
 33463 otherwise, an error number shall be returned to indicate the error.

33464 **ERRORS**

33465 The *pthread\_getschedparam()* function may fail if:

33466 [ESRCH] The value specified by *thread* does not refer to a existing thread.

33467 The *pthread\_setschedparam()* function may fail if:

33468 [EINVAL] The value specified by *policy* or one of the scheduling parameters associated  
33469 with the scheduling policy *policy* is invalid.

33470 [ENOTSUP] An attempt was made to set the policy or scheduling parameters to an  
33471 unsupported value.

33472 TSP [ENOTSUP] An attempt was made to dynamically change the scheduling policy to  
33473 SCHED\_SPORADIC, and the implementation does not support this change.

33474 [EPERM] The caller does not have the appropriate permission to set either the  
33475 scheduling parameters or the scheduling policy of the specified thread.

33476 [EPERM] The implementation does not allow the application to modify one of the  
33477 parameters to the value specified.

33478 [ESRCH] The value specified by *thread* does not refer to a existing thread.

33479 These functions shall not return an error code of [EINTR].

33480 **EXAMPLES**

33481 None.

33482 **APPLICATION USAGE**

33483 None.

33484 **RATIONALE**

33485 None.

33486 **FUTURE DIRECTIONS**

33487 None.

33488 **SEE ALSO**

33489 *pthread\_setschedprio()*, *sched\_getparam()*, *sched\_getscheduler()*, the Base Definitions volume of |  
33490 IEEE Std 1003.1-200x, <pthread.h>, <sched.h>

33491 **CHANGE HISTORY**

33492 First released in Issue 5. Included for alignment with the POSIX Threads Extension.

33493 **Issue 6**

33494 The *pthread\_getschedparam()* and *pthread\_setschedparam()* functions are marked as part of the  
33495 Threads and Thread Execution Scheduling options.

33496 The [ENOSYS] error condition has been removed as stubs need not be provided if an  
33497 implementation does not support the Thread Execution Scheduling option.

33498 The Open Group Corrigendum U026/2 is applied, correcting the prototype for the  
33499 *pthread\_setschedparam()* function so that its second argument is of type **int**.

33500 The SCHED\_SPORADIC scheduling policy is added for alignment with IEEE Std 1003.1d-1999.

33501 The **restrict** keyword is added to the *pthread\_getschedparam()* prototype for alignment with the  
33502 ISO/IEC 9899:1999 standard.

33503 The Open Group Corrigendum U047/1 is applied. |

33504  
33505

IEEE PASC Interpretation 1003.1 #96 is applied, noting that priority values can also be set by a call to the *pthread\_setschedprio()* function.



33506 **NAME**

33507 pthread\_getspecific, pthread\_setspecific — thread-specific data management

33508 **SYNOPSIS**

33509 THR #include &lt;pthread.h&gt;

33510 void \*pthread\_getspecific(pthread\_key\_t key);

33511 int pthread\_setspecific(pthread\_key\_t key, const void \*value);

33512

33513 **DESCRIPTION**33514 The *pthread\_getspecific()* function shall return the value currently bound to the specified *key* on  
33515 behalf of the calling thread.33516 The *pthread\_setspecific()* function shall associate a thread-specific *value* with a *key* obtained via a  
33517 previous call to *pthread\_key\_create()*. Different threads may bind different values to the same  
33518 key. These values are typically pointers to blocks of dynamically allocated memory that have  
33519 been reserved for use by the calling thread.33520 The effect of calling *pthread\_getspecific()* or *pthread\_setspecific()* with a *key* value not obtained  
33521 from *pthread\_key\_create()* or after *key* has been deleted with *pthread\_key\_delete()* is undefined.33522 Both *pthread\_getspecific()* and *pthread\_setspecific()* may be called from a thread-specific data  
33523 destructor function. A call to *pthread\_getspecific()* for the thread-specific data key being  
33524 destroyed shall return the value NULL, unless the value is changed (after the destructor starts)  
33525 by a call to *pthread\_setspecific()*. Calling *pthread\_setspecific()* from a thread-specific data  
33526 destructor routine may result either in lost storage (after at least  
33527 PTHREAD\_DESTRUCTOR\_ITERATIONS attempts at destruction) or in an infinite loop.

33528 Both functions may be implemented as macros.

33529 **RETURN VALUE**33530 The *pthread\_getspecific()* function shall return the thread-specific data value associated with the  
33531 given *key*. If no thread-specific data value is associated with *key*, then the value NULL shall be  
33532 returned.33533 If successful, the *pthread\_setspecific()* function shall return zero; otherwise, an error number shall  
33534 be returned to indicate the error.33535 **ERRORS**33536 No errors are returned from *pthread\_getspecific()*.33537 The *pthread\_setspecific()* function shall fail if:

33538 [ENOMEM] Insufficient memory exists to associate the value with the key.

33539 The *pthread\_setspecific()* function may fail if:

33540 [EINVAL] The key value is invalid.

33541 These functions shall not return an error code of [EINTR].

33542 **EXAMPLES**

33543           None.

33544 **APPLICATION USAGE**

33545           None.

33546 **RATIONALE**

33547           Performance and ease-of-use of *pthread\_getspecific()* is critical for functions that rely on  
33548           maintaining state in thread-specific data. Since no errors are required to be detected by it, and  
33549           since the only error that could be detected is the use of an invalid key, the function to  
33550           *pthread\_getspecific()* has been designed to favor speed and simplicity over error reporting.

33551 **FUTURE DIRECTIONS**

33552           None.

33553 **SEE ALSO**

33554           *pthread\_key\_create()*, the Base Definitions volume of IEEE Std 1003.1-200x, <pthread.h>

33555 **CHANGE HISTORY**

33556           First released in Issue 5. Included for alignment with the POSIX Threads Extension.

33557 **Issue 6**

33558           The *pthread\_getspecific()* and *pthread\_setspecific()* functions are marked as part of the Threads  
33559           option.

33560           IEEE PASC Interpretation 1003.1c #3 (Part 6) is applied, updating the DESCRIPTION.

33561 **NAME**

33562 pthread\_join — wait for thread termination

33563 **SYNOPSIS**

33564 THR #include &lt;pthread.h&gt;

33565 int pthread\_join(pthread\_t thread, void \*\*value\_ptr);

33566

33567 **DESCRIPTION**

33568 The *pthread\_join()* function shall suspend execution of the calling thread until the target *thread*  
 33569 terminates, unless the target *thread* has already terminated. On return from a successful  
 33570 *pthread\_join()* call with a non-NULL *value\_ptr* argument, the value passed to *pthread\_exit()* by  
 33571 the terminating thread shall be made available in the location referenced by *value\_ptr*. When a  
 33572 *pthread\_join()* returns successfully, the target thread has been terminated. The results of multiple  
 33573 simultaneous calls to *pthread\_join()* specifying the same target thread are undefined. If the  
 33574 thread calling *pthread\_join()* is canceled, then the target thread shall not be detached.

33575 It is unspecified whether a thread that has exited but remains unjoined counts against  
 33576 \_PTHREAD\_THREADS\_MAX.

33577 **RETURN VALUE**

33578 If successful, the *pthread\_join()* function shall return zero; otherwise, an error number shall be  
 33579 returned to indicate the error.

33580 **ERRORS**33581 The *pthread\_join()* function shall fail if:

33582 [EINVAL] The implementation has detected that the value specified by *thread* does not  
 33583 refer to a joinable thread.

33584 [ESRCH] No thread could be found corresponding to that specified by the given thread  
 33585 ID.

33586 The *pthread\_join()* function may fail if:

33587 [EDEADLK] A deadlock was detected or the value of *thread* specifies the calling thread.

33588 The *pthread\_join()* function shall not return an error code of [EINTR].33589 **EXAMPLES**

33590 An example of thread creation and deletion follows:

```

33591 typedef struct {
33592     int *ar;
33593     long n;
33594 } subarray;
33595
33596 void *
33597 incer(void *arg)
33598 {
33599     long i;
33599     for (i = 0; i < ((subarray *)arg)->n; i++)
33600         ((subarray *)arg)->ar[i]++;
33601 }
33602
33603 main()
33604 {
33604     int ar[1000000];

```

```

33605     pthread_t  th1, th2;
33606     subarray  sb1, sb2;

33607     sb1.ar = &ar[0];
33608     sb1.n  = 500000;
33609     (void) pthread_create(&th1, NULL, incer, &sb1);

33610     sb2.ar = &ar[500000];
33611     sb2.n  = 500000;
33612     (void) pthread_create(&th2, NULL, incer, &sb2);

33613     (void) pthread_join(th1, NULL);
33614     (void) pthread_join(th2, NULL);
33615 }

```

**33616 APPLICATION USAGE**

33617 None.

**33618 RATIONALE**

33619 The *pthread\_join()* function is a convenience that has proven useful in multi-threaded  
33620 applications. It is true that a programmer could simulate this function if it were not provided by  
33621 passing extra state as part of the argument to the *start\_routine()*. The terminating thread would  
33622 set a flag to indicate termination and broadcast a condition that is part of that state; a joining  
33623 thread would wait on that condition variable. While such a technique would allow a thread to  
33624 wait on more complex conditions (for example, waiting for multiple threads to terminate),  
33625 waiting on individual thread termination is considered widely useful. Also, including the  
33626 *pthread\_join()* function in no way precludes a programmer from coding such complex waits.  
33627 Thus, while not a primitive, including *pthread\_join()* in this volume of IEEE Std 1003.1-200x was  
33628 considered valuable.

33629 The *pthread\_join()* function provides a simple mechanism allowing an application to wait for a  
33630 thread to terminate. After the thread terminates, the application may then choose to clean up  
33631 resources that were used by the thread. For instance, after *pthread\_join()* returns, any  
33632 application-provided stack storage could be reclaimed.

33633 The *pthread\_join()* or *pthread\_detach()* function should eventually be called for every thread that  
33634 is created with the *detachstate* attribute set to *PTHREAD\_CREATE\_JOINABLE* so that storage  
33635 associated with the thread may be reclaimed.

33636 The interaction between *pthread\_join()* and cancelation is well-defined for the following reasons:

- 33637 • The *pthread\_join()* function, like all other non-async-cancel-safe functions, can only be called  
33638 with deferred cancelability type.
- 33639 • Cancelation cannot occur in the disabled cancelability state.

33640 Thus, only the default cancelability state need be considered. As specified, either the  
33641 *pthread\_join()* call is canceled, or it succeeds, but not both. The difference is obvious to the  
33642 application, since either a cancelation handler is run or *pthread\_join()* returns. There are no race  
33643 conditions since *pthread\_join()* was called in the deferred cancelability state.

**33644 FUTURE DIRECTIONS**

33645 None.

**33646 SEE ALSO**

33647 *pthread\_create()*, *wait()*, the Base Definitions volume of IEEE Std 1003.1-200x, <**pthread.h**>

33648 **CHANGE HISTORY**

33649 First released in Issue 5. Included for alignment with the POSIX Threads Extension.

33650 **Issue 6**

33651 The *pthread\_join()* function is marked as part of the Threads option.

33652 **NAME**

33653 pthread\_key\_create — thread-specific data key creation

33654 **SYNOPSIS**

33655 THR #include &lt;pthread.h&gt;

33656 int pthread\_key\_create(pthread\_key\_t \*key, void (\*destructor)(void\*));

33657

33658 **DESCRIPTION**

33659 The *pthread\_key\_create()* function shall create a thread-specific data key visible to all threads in  
33660 the process. Key values provided by *pthread\_key\_create()* are opaque objects used to locate  
33661 thread-specific data. Although the same key value may be used by different threads, the values  
33662 bound to the key by *pthread\_setspecific()* are maintained on a per-thread basis and persist for the  
33663 life of the calling thread.

33664 Upon key creation, the value NULL shall be associated with the new key in all active threads.  
33665 Upon thread creation, the value NULL shall be associated with all defined keys in the new  
33666 thread.

33667 An optional destructor function may be associated with each key value. At thread exit, if a key  
33668 value has a non-NULL destructor pointer, and the thread has a non-NULL value associated with  
33669 that key, the value of the key is set to NULL, and then the function pointed to is called with the  
33670 previously associated value as its sole argument. The order of destructor calls is unspecified if  
33671 more than one destructor exists for a thread when it exits.

33672 If, after all the destructors have been called for all non-NULL values with associated destructors,  
33673 there are still some non-NULL values with associated destructors, then the process is repeated.  
33674 If, after at least {PTHREAD\_DESTRUCTOR\_ITERATIONS} iterations of destructor calls for  
33675 outstanding non-NULL values, there are still some non-NULL values with associated  
33676 destructors, implementations may stop calling destructors, or they may continue calling  
33677 destructors until no non-NULL values with associated destructors exist, even though this might  
33678 result in an infinite loop.

33679 **RETURN VALUE**

33680 If successful, the *pthread\_key\_create()* function shall store the newly created key value at *\*key* and  
33681 shall return zero. Otherwise, an error number shall be returned to indicate the error.

33682 **ERRORS**33683 The *pthread\_key\_create()* function shall fail if:

33684 [EAGAIN] The system lacked the necessary resources to create another thread-specific  
33685 data key, or the system-imposed limit on the total number of keys per process  
33686 PTHREAD\_KEYS\_MAX has been exceeded.

33687 [ENOMEM] Insufficient memory exists to create the key.

33688 The *pthread\_key\_create()* function shall not return an error code of [EINTR].

33689 **EXAMPLES**

33690 The following example demonstrates a function that initializes a thread-specific data key when  
 33691 it is first called, and associates a thread-specific object with each calling thread, initializing this  
 33692 object when necessary.

```

33693     static pthread_key_t key;
33694     static pthread_once_t key_once = PTHREAD_ONCE_INIT;

33695     static void
33696     make_key()
33697     {
33698         (void) pthread_key_create(&key, NULL);
33699     }

33700     func()
33701     {
33702         void *ptr;

33703         (void) pthread_once(&key_once, make_key);
33704         if ((ptr = pthread_getspecific(key)) == NULL) {
33705             ptr = malloc(OBJECT_SIZE);
33706             ...
33707             (void) pthread_setspecific(key, ptr);
33708         }
33709         ...
33710     }
  
```

33711 Note that the key has to be initialized before *pthread\_getspecific()* or *pthread\_setspecific()* can be  
 33712 used. The *pthread\_key\_create()* call could either be explicitly made in a module initialization  
 33713 routine, or it can be done implicitly by the first call to a module as in this example. Any attempt  
 33714 to use the key before it is initialized is a programming error, making the code below incorrect.

```

33715     static pthread_key_t key;

33716     func()
33717     {
33718         void *ptr;

33719         /* KEY NOT INITIALIZED!!! THIS WON'T WORK!!! */
33720         if ((ptr = pthread_getspecific(key)) == NULL &&
33721             pthread_setspecific(key, NULL) != 0) {
33722             pthread_key_create(&key, NULL);
33723             ...
33724         }
33725     }
  
```

33726 **APPLICATION USAGE**

33727 None.

## 33728 RATIONALE

33729 **Destructor Functions**

33730 Normally, the value bound to a key on behalf of a particular thread is a pointer to storage  
33731 allocated dynamically on behalf of the calling thread. The destructor functions specified with  
33732 *pthread\_key\_create()* are intended to be used to free this storage when the thread exits. Thread  
33733 cancellation cleanup handlers cannot be used for this purpose because thread-specific data may  
33734 persist outside the lexical scope in which the cancellation cleanup handlers operate.

33735 If the value associated with a key needs to be updated during the lifetime of the thread, it may  
33736 be necessary to release the storage associated with the old value before the new value is bound.  
33737 Although the *pthread\_setspecific()* function could do this automatically, this feature is not needed  
33738 often enough to justify the added complexity. Instead, the programmer is responsible for freeing  
33739 the stale storage:

```
33740 pthread_getspecific(key, &old);  
33741 new = allocate();  
33742 destructor(old);  
33743 pthread_setspecific(key, new);
```

33744 **Note:** The above example could leak storage if run with asynchronous cancellation enabled. No such  
33745 problems occur in the default cancellation state if no cancellation points occur between the get  
33746 and set.

33747 There is no notion of a destructor-safe function. If an application does not call *pthread\_exit()*  
33748 from a signal handler, or if it blocks any signal whose handler may call *pthread\_exit()* while  
33749 calling async-unsafe functions, all functions may be safely called from destructors.

33750 **Non-Idempotent Data Key Creation**

33751 There were requests to make *pthread\_key\_create()* idempotent with respect to a given *key* address  
33752 parameter. This would allow applications to call *pthread\_key\_create()* multiple times for a given  
33753 *key* address and be guaranteed that only one key would be created. Doing so would require the  
33754 key value to be previously initialized (possibly at compile time) to a known null value and  
33755 would require that implicit mutual-exclusion be performed based on the address and contents of  
33756 the *key* parameter in order to guarantee that exactly one key would be created.

33757 Unfortunately, the implicit mutual-exclusion would not be limited to only *pthread\_key\_create()*.  
33758 On many implementations, implicit mutual-exclusion would also have to be performed by  
33759 *pthread\_getspecific()* and *pthread\_setspecific()* in order to guard against using incompletely stored  
33760 or not-yet-visible key values. This could significantly increase the cost of important operations,  
33761 particularly *pthread\_getspecific()*.

33762 Thus, this proposal was rejected. The *pthread\_key\_create()* function performs no implicit  
33763 synchronization. It is the responsibility of the programmer to ensure that it is called exactly once  
33764 per key before use of the key. Several straightforward mechanisms can already be used to  
33765 accomplish this, including calling explicit module initialization functions, using mutexes, and  
33766 using *pthread\_once()*. This places no significant burden on the programmer, introduces no  
33767 possibly confusing *ad hoc* implicit synchronization mechanism, and potentially allows  
33768 commonly used thread-specific data operations to be more efficient.

33769 **FUTURE DIRECTIONS**

33770 None.



33771 **SEE ALSO**

33772 *pthread\_getspecific()*, *pthread\_key\_delete()*, the Base Definitions volume of IEEE Std 1003.1-200x,  
33773 <pthread.h>

33774 **CHANGE HISTORY**

33775 First released in Issue 5. Included for alignment with the POSIX Threads Extension.

33776 **Issue 6**

33777 The *pthread\_key\_create()* function is marked as part of the Threads option.

33778 IEEE PASC Interpretation 1003.1c #8 is applied, updating the DESCRIPTION.

33779 **NAME**

33780 pthread\_key\_delete — thread-specific data key deletion

33781 **SYNOPSIS**

33782 THR #include &lt;pthread.h&gt;

33783 int pthread\_key\_delete(pthread\_key\_t key);

33784

33785 **DESCRIPTION**

33786 The *pthread\_key\_delete()* function shall delete a thread-specific data key previously returned by  
33787 *pthread\_key\_create()*. The thread-specific data values associated with *key* need not be NULL at  
33788 the time *pthread\_key\_delete()* is called. It is the responsibility of the application to free any  
33789 application storage or perform any cleanup actions for data structures related to the deleted key  
33790 or associated thread-specific data in any threads; this cleanup can be done either before or after  
33791 *pthread\_key\_delete()* is called. Any attempt to use *key* following the call to *pthread\_key\_delete()*  
33792 results in undefined behavior.

33793 The *pthread\_key\_delete()* function shall be callable from within destructor functions. No  
33794 destructor functions shall be invoked by *pthread\_key\_delete()*. Any destructor function that may  
33795 have been associated with *key* shall no longer be called upon thread exit.

33796 **RETURN VALUE**

33797 If successful, the *pthread\_key\_delete()* function shall return zero; otherwise, an error number shall  
33798 be returned to indicate the error.

33799 **ERRORS**33800 The *pthread\_key\_delete()* function may fail if:33801 [EINVAL] The *key* value is invalid.33802 The *pthread\_key\_delete()* function shall not return an error code of [EINTR].33803 **EXAMPLES**

33804 None.

33805 **APPLICATION USAGE**

33806 None.

33807 **RATIONALE**

33808 A thread-specific data key deletion function has been included in order to allow the resources  
33809 associated with an unused thread-specific data key to be freed. Unused thread-specific data keys  
33810 can arise, among other scenarios, when a dynamically loaded module that allocated a key is  
33811 unloaded.

33812 Conforming applications are responsible for performing any cleanup actions needed for data  
33813 structures associated with the key to be deleted, including data referenced by thread-specific  
33814 data values. No such cleanup is done by *pthread\_key\_delete()*. In particular, destructor functions  
33815 are not called. There are several reasons for this division of responsibility:

- 33816 1. The associated destructor functions used to free thread-specific data at thread exit time are  
33817 only guaranteed to work correctly when called in the thread that allocated the thread-  
33818 specific data. (Destructors themselves may utilize thread-specific data.) Thus, they cannot  
33819 be used to free thread-specific data in other threads at key deletion time. Attempting to  
33820 have them called by other threads at key deletion time would require other threads to be  
33821 asynchronously interrupted. But since interrupted threads could be in an arbitrary state,  
33822 including holding locks necessary for the destructor to run, this approach would fail. In  
33823 general, there is no safe mechanism whereby an implementation could free thread-specific  
33824 data at key deletion time.

33825           2. Even if there were a means of safely freeing thread-specific data associated with keys to be  
33826 deleted, doing so would require that implementations be able to enumerate the threads  
33827 with non-NULL data and potentially keep them from creating more thread-specific data  
33828 while the key deletion is occurring. This special case could cause extra synchronization in  
33829 the normal case, which would otherwise be unnecessary.

33830           For an application to know that it is safe to delete a key, it has to know that all the threads that  
33831 might potentially ever use the key do not attempt to use it again. For example, it could know this  
33832 if all the client threads have called a cleanup procedure declaring that they are through with the  
33833 module that is being shut down, perhaps by zero'ing a reference count.

33834 **FUTURE DIRECTIONS**

33835           None.

33836 **SEE ALSO**

33837           *pthread\_key\_create()*, the Base Definitions volume of IEEE Std 1003.1-200x, <pthread.h>

33838 **CHANGE HISTORY**

33839           First released in Issue 5. Included for alignment with the POSIX Threads Extension.

33840 **Issue 6**

33841           The *pthread\_key\_delete()* function is marked as part of the Threads option.

33842 **NAME**

33843 pthread\_kill — send a signal to a thread

33844 **SYNOPSIS**

33845 THR #include &lt;signal.h&gt;

33846 int pthread\_kill(pthread\_t thread, int sig);

33847

33848 **DESCRIPTION**33849 The *pthread\_kill()* function shall request that a signal be delivered to the specified thread. |33850 As in *kill()*, if *sig* is zero, error checking shall be performed but no signal shall actually be sent. |33851 **RETURN VALUE**

33852 Upon successful completion, the function shall return a value of zero. Otherwise, the function

33853 shall return an error number. If the *pthread\_kill()* function fails, no signal shall be sent.33854 **ERRORS**33855 The *pthread\_kill()* function shall fail if:33856 [ESRCH] No thread could be found corresponding to that specified by the given thread  
33857 ID.33858 [EINVAL] The value of the *sig* argument is an invalid or unsupported signal number.33859 The *pthread\_kill()* function shall not return an error code of [EINTR].33860 **EXAMPLES**

33861 None.

33862 **APPLICATION USAGE**33863 The *pthread\_kill()* function provides a mechanism for asynchronously directing a signal at a  
33864 thread in the calling process. This could be used, for example, by one thread to affect broadcast  
33865 delivery of a signal to a set of threads.33866 Note that *pthread\_kill()* only causes the signal to be handled in the context of the given thread;  
33867 the signal action (termination or stopping) affects the process as a whole.33868 **RATIONALE**

33869 None.

33870 **FUTURE DIRECTIONS**

33871 None.

33872 **SEE ALSO**33873 *kill()*, *pthread\_self()*, *raise()*, the Base Definitions volume of IEEE Std 1003.1-200x, <signal.h>33874 **CHANGE HISTORY**

33875 First released in Issue 5. Included for alignment with the POSIX Threads Extension.

33876 **Issue 6**33877 The *pthread\_kill()* function is marked as part of the Threads option.

33878 The APPLICATION USAGE section is added.

33879 **NAME**

33880 pthread\_mutex\_destroy, pthread\_mutex\_init — destroy and initialize a mutex

33881 **SYNOPSIS**

33882 THR #include &lt;pthread.h&gt;

33883 int pthread\_mutex\_destroy(pthread\_mutex\_t \*mutex);

33884 int pthread\_mutex\_init(pthread\_mutex\_t \*restrict mutex,

33885 const pthread\_mutexattr\_t \*restrict attr);

33886 pthread\_mutex\_t mutex = PTHREAD\_MUTEX\_INITIALIZER;

33887

33888 **DESCRIPTION**

33889 The *pthread\_mutex\_destroy()* function shall destroy the mutex object referenced by *mutex*; the  
 33890 mutex object becomes, in effect, uninitialized. An implementation may cause  
 33891 *pthread\_mutex\_destroy()* to set the object referenced by *mutex* to an invalid value. A destroyed  
 33892 mutex object can be reinitialized using *pthread\_mutex\_init()*; the results of otherwise referencing  
 33893 the object after it has been destroyed are undefined.

33894 It shall be safe to destroy an initialized mutex that is unlocked. Attempting to destroy a locked  
 33895 mutex results in undefined behavior.

33896 The *pthread\_mutex\_init()* function shall initialize the mutex referenced by *mutex* with attributes  
 33897 specified by *attr*. If *attr* is NULL, the default mutex attributes are used; the effect shall be the  
 33898 same as passing the address of a default mutex attributes object. Upon successful initialization,  
 33899 the state of the mutex becomes initialized and unlocked.

33900 Only *mutex* itself may be used for performing synchronization. The result of referring to copies  
 33901 of *mutex* in calls to *pthread\_mutex\_lock()*, *pthread\_mutex\_trylock()*, *pthread\_mutex\_unlock()*, and  
 33902 *pthread\_mutex\_destroy()* is undefined.

33903 Attempting to initialize an already initialized mutex results in undefined behavior.

33904 In cases where default mutex attributes are appropriate, the macro  
 33905 PTHREAD\_MUTEX\_INITIALIZER can be used to initialize mutexes that are statically allocated.  
 33906 The effect shall be equivalent to dynamic initialization by a call to *pthread\_mutex\_init()* with  
 33907 parameter *attr* specified as NULL, except that no error checks are performed.

33908 **RETURN VALUE**

33909 If successful, the *pthread\_mutex\_destroy()* and *pthread\_mutex\_init()* functions shall return zero;  
 33910 otherwise, an error number shall be returned to indicate the error.

33911 The [EBUSY] and [EINVAL] error checks, if implemented, act as if they were performed  
 33912 immediately at the beginning of processing for the function and shall cause an error return prior  
 33913 to modifying the state of the mutex specified by *mutex*.

33914 **ERRORS**

33915 The *pthread\_mutex\_destroy()* function may fail if:

33916 [EBUSY] The implementation has detected an attempt to destroy the object referenced  
 33917 by *mutex* while it is locked or referenced (for example, while being used in a  
 33918 *pthread\_cond\_timedwait()* or *pthread\_cond\_wait()*) by another thread.

33919 [EINVAL] The value specified by *mutex* is invalid.

33920 The *pthread\_mutex\_init()* function shall fail if:

33921 [EAGAIN] The system lacked the necessary resources (other than memory) to initialize  
 33922 another mutex.

- 33923 [ENOMEM] Insufficient memory exists to initialize the mutex.
- 33924 [EPERM] The caller does not have the privilege to perform the operation.
- 33925 The *pthread\_mutex\_init()* function may fail if:
- 33926 [EBUSY] The implementation has detected an attempt to reinitialize the object |  
33927 referenced by *mutex*, a previously initialized, but not yet destroyed, mutex.
- 33928 [EINVAL] The value specified by *attr* is invalid.
- 33929 These functions shall not return an error code of [EINTR].

**33930 EXAMPLES**

33931 None.

**33932 APPLICATION USAGE**

33933 None.

**33934 RATIONALE****33935 Alternate Implementations Possible**

33936 This volume of IEEE Std 1003.1-200x supports several alternative implementations of mutexes.  
33937 An implementation may store the lock directly in the object of type **pthread\_mutex\_t**.  
33938 Alternatively, an implementation may store the lock in the heap and merely store a pointer,  
33939 handle, or unique ID in the mutex object. Either implementation has advantages or may be  
33940 required on certain hardware configurations. So that portable code can be written that is  
33941 invariant to this choice, this volume of IEEE Std 1003.1-200x does not define assignment or  
33942 equality for this type, and it uses the term “initialize” to reinforce the (more restrictive) notion  
33943 that the lock may actually reside in the mutex object itself.

33944 Note that this precludes an over-specification of the type of the mutex or condition variable and  
33945 motivates the opacity of the type.

33946 An implementation is permitted, but not required, to have *pthread\_mutex\_destroy()* store an  
33947 illegal value into the mutex. This may help detect erroneous programs that try to lock (or  
33948 otherwise reference) a mutex that has already been destroyed.

**33949 Tradeoff Between Error Checks and Performance Supported**

33950 Many of the error checks were made optional in order to let implementations trade off  
33951 performance *versus* degree of error checking according to the needs of their specific applications  
33952 and execution environment. As a general rule, errors or conditions caused by the system (such as  
33953 insufficient memory) always need to be reported, but errors due to an erroneously coded  
33954 application (such as failing to provide adequate synchronization to prevent a mutex from being  
33955 deleted while in use) are made optional.

33956 A wide range of implementations is thus made possible. For example, an implementation  
33957 intended for application debugging may implement all of the error checks, but an  
33958 implementation running a single, provably correct application under very tight performance  
33959 constraints in an embedded computer might implement minimal checks. An implementation  
33960 might even be provided in two versions, similar to the options that compilers provide: a full-  
33961 checking, but slower version; and a limited-checking, but faster version. To forbid this  
33962 optionality would be a disservice to users.

33963 By carefully limiting the use of “undefined behavior” only to things that an erroneous (badly  
33964 coded) application might do, and by defining that resource-not-available errors are mandatory,  
33965 this volume of IEEE Std 1003.1-200x ensures that a fully-conforming application is portable

33966 across the full range of implementations, while not forcing all implementations to add overhead  
33967 to check for numerous things that a correct program never does.

### 33968 **Why No Limits Defined**

33969 Defining symbols for the maximum number of mutexes and condition variables was considered  
33970 but rejected because the number of these objects may change dynamically. Furthermore, many  
33971 implementations place these objects into application memory; thus, there is no explicit  
33972 maximum.

### 33973 **Static Initializers for Mutexes and Condition Variables**

33974 Providing for static initialization of statically allocated synchronization objects allows modules  
33975 with private static synchronization variables to avoid runtime initialization tests and overhead.  
33976 Furthermore, it simplifies the coding of self-initializing modules. Such modules are common in  
33977 C libraries, where for various reasons the design calls for self-initialization instead of requiring  
33978 an explicit module initialization function to be called. An example use of static initialization  
33979 follows.

33980 Without static initialization, a self-initializing routine *foo()* might look as follows:

```
33981 static pthread_once_t foo_once = PTHREAD_ONCE_INIT;
33982 static pthread_mutex_t foo_mutex;

33983 void foo_init()
33984 {
33985     pthread_mutex_init(&foo_mutex, NULL);
33986 }

33987 void foo()
33988 {
33989     pthread_once(&foo_once, foo_init);
33990     pthread_mutex_lock(&foo_mutex);
33991     /* Do work. */
33992     pthread_mutex_unlock(&foo_mutex);
33993 }
```

33994 With static initialization, the same routine could be coded as follows:

```
33995 static pthread_mutex_t foo_mutex = PTHREAD_MUTEX_INITIALIZER;

33996 void foo()
33997 {
33998     pthread_mutex_lock(&foo_mutex);
33999     /* Do work. */
34000     pthread_mutex_unlock(&foo_mutex);
34001 }
```

34002 Note that the static initialization both eliminates the need for the initialization test inside  
34003 *pthread\_once()* and the fetch of *&foo\_mutex* to learn the address to be passed to  
34004 *pthread\_mutex\_lock()* or *pthread\_mutex\_unlock()*.

34005 Thus, the C code written to initialize static objects is simpler on all systems and is also faster on a  
34006 large class of systems; those where the (entire) synchronization object can be stored in  
34007 application memory.

34008 Yet the locking performance question is likely to be raised for machines that require mutexes to  
34009 be allocated out of special memory. Such machines actually have to have mutexes and possibly

34010 condition variables contain pointers to the actual hardware locks. For static initialization to work  
34011 on such machines, *pthread\_mutex\_lock()* also has to test whether or not the pointer to the actual  
34012 lock has been allocated. If it has not, *pthread\_mutex\_lock()* has to initialize it before use. The  
34013 reservation of such resources can be made when the program is loaded, and hence return codes  
34014 have not been added to mutex locking and condition variable waiting to indicate failure to  
34015 complete initialization.

34016 This runtime test in *pthread\_mutex\_lock()* would at first seem to be extra work; an extra test is  
34017 required to see whether the pointer has been initialized. On most machines this would actually  
34018 be implemented as a fetch of the pointer, testing the pointer against zero, and then using the  
34019 pointer if it has already been initialized. While the test might seem to add extra work, the extra  
34020 effort of testing a register is usually negligible since no extra memory references are actually  
34021 done. As more and more machines provide caches, the real expenses are memory references, not  
34022 instructions executed.

34023 Alternatively, depending on the machine architecture, there are often ways to eliminate *all*  
34024 overhead in the most important case: on the lock operations that occur *after* the lock has been  
34025 initialized. This can be done by shifting more overhead to the less frequent operation:  
34026 initialization. Since out-of-line mutex allocation also means that an address has to be  
34027 dereferenced to find the actual lock, one technique that is widely applicable is to have static  
34028 initialization store a bogus value for that address; in particular, an address that causes a machine  
34029 fault to occur. When such a fault occurs upon the first attempt to lock such a mutex, validity  
34030 checks can be done, and then the correct address for the actual lock can be filled in. Subsequent  
34031 lock operations incur no extra overhead since they do not “fault”. This is merely one technique  
34032 that can be used to support static initialization, while not adversely affecting the performance of  
34033 lock acquisition. No doubt there are other techniques that are highly machine-dependent.

34034 The locking overhead for machines doing out-of-line mutex allocation is thus similar for  
34035 modules being implicitly initialized, where it is improved for those doing mutex allocation  
34036 entirely inline. The inline case is thus made much faster, and the out-of-line case is not  
34037 significantly worse.

34038 Besides the issue of locking performance for such machines, a concern is raised that it is possible  
34039 that threads would serialize contending for initialization locks when attempting to finish  
34040 initializing statically allocated mutexes. (Such finishing would typically involve taking an  
34041 internal lock, allocating a structure, storing a pointer to the structure in the mutex, and releasing  
34042 the internal lock.) First, many implementations would reduce such serialization by hashing on  
34043 the mutex address. Second, such serialization can only occur a bounded number of times. In  
34044 particular, it can happen at most as many times as there are statically allocated synchronization  
34045 objects. Dynamically allocated objects would still be initialized via *pthread\_mutex\_init()* or  
34046 *pthread\_cond\_init()*.

34047 Finally, if none of the above optimization techniques for out-of-line allocation yields sufficient  
34048 performance for an application on some implementation, the application can avoid static  
34049 initialization altogether by explicitly initializing all synchronization objects with the  
34050 corresponding *pthread\_\*\_init()* functions, which are supported by all implementations. An  
34051 implementation can also document the tradeoffs and advise which initialization technique is  
34052 more efficient for that particular implementation.



34053 **Destroying Mutexes**

34054 A mutex can be destroyed immediately after it is unlocked. For example, consider the following  
 34055 code:

```

34056 struct obj {
34057     pthread_mutex_t om;
34058     int refcnt;
34059     ...
34060 };

34061 obj_done(struct obj *op)
34062 {
34063     pthread_mutex_lock(&op->om);
34064     if (--op->refcnt == 0) {
34065         pthread_mutex_unlock(&op->om);
34066         (A) pthread_mutex_destroy(&op->om);
34067         (B) free(op);
34068     } else
34069         (C) pthread_mutex_unlock(&op->om);
34070 }

```

34071 In this case *obj* is reference counted and *obj\_done()* is called whenever a reference to the object is  
 34072 dropped. Implementations are required to allow an object to be destroyed and freed and  
 34073 potentially unmapped (for example, lines A and B) immediately after the object is unlocked (line  
 34074 C).

34075 **FUTURE DIRECTIONS**

34076 None.

34077 **SEE ALSO**

34078 *pthread\_mutex\_getprioceiling()*, *pthread\_mutex\_lock()*, *pthread\_mutex\_timedlock()*,  
 34079 *pthread\_mutexattr\_getpshared()*, the Base Definitions volume of IEEE Std 1003.1-200x,  
 34080 <pthread.h>

34081 **CHANGE HISTORY**

34082 First released in Issue 5. Included for alignment with the POSIX Threads Extension.

34083 **Issue 6**

34084 The *pthread\_mutex\_destroy()* and *pthread\_mutex\_init()* functions are marked as part of the  
 34085 Threads option.

34086 The *pthread\_mutex\_timedlock()* function is added to the SEE ALSO section for alignment with  
 34087 IEEE Std 1003.1d-1999.

34088 IEEE PASC Interpretation 1003.1c #34 is applied, updating the DESCRIPTION.

34089 The **restrict** keyword is added to the *pthread\_mutex\_init()* prototype for alignment with the  
 34090 ISO/IEC 9899:1999 standard.

## 34091 NAME

34092 pthread\_mutex\_getprioceiling, pthread\_mutex\_setprioceiling — get and set the priority ceiling  
 34093 of a mutex (**REALTIME THREADS**)

## 34094 SYNOPSIS

```
34095 THR TPP #include <pthread.h>
34096
34096 int pthread_mutex_getprioceiling(const pthread_mutex_t *restrict mutex,
34097 int *restrict prioceiling);
34098 int pthread_mutex_setprioceiling(pthread_mutex_t *restrict mutex,
34099 int prioceiling, int *restrict old_ceiling);
34100
```

## 34101 DESCRIPTION

34102 The *pthread\_mutex\_getprioceiling()* function shall return the current priority ceiling of the mutex.

34103 The *pthread\_mutex\_setprioceiling()* function shall either lock the mutex if it is unlocked, or block |  
 34104 until it can successfully lock the mutex, then it shall change the mutex's priority ceiling and |  
 34105 release the mutex. When the change is successful, the previous value of the priority ceiling shall |  
 34106 be returned in *old\_ceiling*. The process of locking the mutex need not adhere to the priority |  
 34107 protect protocol.

34108 If the *pthread\_mutex\_setprioceiling()* function fails, the mutex priority ceiling shall not be  
 34109 changed.

## 34110 RETURN VALUE

34111 If successful, the *pthread\_mutex\_getprioceiling()* and *pthread\_mutex\_setprioceiling()* functions shall  
 34112 return zero; otherwise, an error number shall be returned to indicate the error.

## 34113 ERRORS

34114 The *pthread\_mutex\_getprioceiling()* and *pthread\_mutex\_setprioceiling()* functions may fail if:

- 34115 [EINVAL] The priority requested by *prioceiling* is out of range.
- 34116 [EINVAL] The value specified by *mutex* does not refer to a currently existing mutex.
- 34117 [EPERM] The caller does not have the privilege to perform the operation.
- 34118 These functions shall not return an error code of [EINTR].

## 34119 EXAMPLES

34120 None.

## 34121 APPLICATION USAGE

34122 None.

## 34123 RATIONALE

34124 None.

## 34125 FUTURE DIRECTIONS

34126 None.

## 34127 SEE ALSO

34128 *pthread\_mutex\_destroy()*, *pthread\_mutex\_lock()*, *pthread\_mutex\_timedlock()*, the Base Definitions  
 34129 volume of IEEE Std 1003.1-200x, <pthread.h>

## 34130 CHANGE HISTORY

- 34131 First released in Issue 5. Included for alignment with the POSIX Threads Extension.
- 34132 Marked as part of the Realtime Threads Feature Group.

34133 **Issue 6**

34134 The *pthread\_mutex\_getprioceiling()* and *pthread\_mutex\_setprioceiling()* functions are marked as  
34135 part of the Threads and Thread Priority Protection options.

34136 The [ENOSYS] error condition has been removed as stubs need not be provided if an  
34137 implementation does not support the Thread Priority Protection option.

34138 The [ENOSYS] error denoting non-support of the priority ceiling protocol for mutexes has been  
34139 removed. This is since if the implementation provides the functions (regardless of whether  
34140 `_POSIX_PTHREAD_PRIO_PROTECT` is defined), they must function as in the DESCRIPTION  
34141 and therefore the priority ceiling protocol for mutexes is supported.

34142 The *pthread\_mutex\_timedlock()* function is added to the SEE ALSO section for alignment with  
34143 IEEE Std 1003.1d-1999.

34144 The **restrict** keyword is added to the *pthread\_mutex\_getprioceiling()* and  
34145 *pthread\_mutex\_setprioceiling()* prototypes for alignment with the ISO/IEC 9899:1999 standard.

34146 **NAME**

34147 pthread\_mutex\_init — initialize a mutex

34148 **SYNOPSIS**

34149 THR #include <pthread.h>

34150 int pthread\_mutex\_init(pthread\_mutex\_t \*restrict mutex,

34151 const pthread\_mutexattr\_t \*restrict attr);

34152 pthread\_mutex\_t mutex = PTHREAD\_MUTEX\_INITIALIZER;

34153

34154 **DESCRIPTION**

34155 Refer to *pthread\_mutex\_destroy()*.

34156 **NAME**

34157 pthread\_mutex\_lock, pthread\_mutex\_trylock, pthread\_mutex\_unlock — lock and unlock a  
 34158 mutex

34159 **SYNOPSIS**

34160 THR #include <pthread.h>

34161 int pthread\_mutex\_lock(pthread\_mutex\_t \*mutex);

34162 int pthread\_mutex\_trylock(pthread\_mutex\_t \*mutex);

34163 int pthread\_mutex\_unlock(pthread\_mutex\_t \*mutex);

34164

34165 **DESCRIPTION**

34166 The mutex object referenced by *mutex* shall be locked by calling *pthread\_mutex\_lock()*. If the  
 34167 mutex is already locked, the calling thread shall block until the mutex becomes available. This  
 34168 operation shall return with the mutex object referenced by *mutex* in the locked state with the  
 34169 calling thread as its owner.

34170 XSI If the mutex type is PTHREAD\_MUTEX\_NORMAL, deadlock detection shall not be provided.  
 34171 Attempting to relock the mutex causes deadlock. If a thread attempts to unlock a mutex that it  
 34172 has not locked or a mutex which is unlocked, undefined behavior results.

34173 If the mutex type is PTHREAD\_MUTEX\_ERRORCHECK, then error checking shall be provided.  
 34174 If a thread attempts to relock a mutex that it has already locked, an error shall be returned. If a  
 34175 thread attempts to unlock a mutex that it has not locked or a mutex which is unlocked, an error  
 34176 shall be returned.

34177 If the mutex type is PTHREAD\_MUTEX\_RECURSIVE, then the mutex shall maintain the  
 34178 concept of a lock count. When a thread successfully acquires a mutex for the first time, the lock  
 34179 count shall be set to one. Every time a thread relocks this mutex, the lock count shall be  
 34180 incremented by one. Each time the thread unlocks the mutex, the lock count shall be  
 34181 decremented by one. When the lock count reaches zero, the mutex shall become available for  
 34182 other threads to acquire. If a thread attempts to unlock a mutex that it has not locked or a mutex  
 34183 which is unlocked, an error shall be returned.

34184 If the mutex type is PTHREAD\_MUTEX\_DEFAULT, attempting to recursively lock the mutex  
 34185 results in undefined behavior. Attempting to unlock the mutex if it was not locked by the calling  
 34186 thread results in undefined behavior. Attempting to unlock the mutex if it is not locked results in  
 34187 undefined behavior.

34188 The *pthread\_mutex\_trylock()* function shall be equivalent to *pthread\_mutex\_lock()*, except that if  
 34189 the mutex object referenced by *mutex* is currently locked (by any thread, including the current  
 34190 thread), the call shall return immediately. If the mutex type is PTHREAD\_MUTEX\_RECURSIVE  
 34191 and the mutex is currently owned by the calling thread, the mutex lock count shall be  
 34192 incremented by one and the *pthread\_mutex\_trylock()* function shall immediately return success.

34193 XSI The *pthread\_mutex\_unlock()* function shall release the mutex object referenced by *mutex*. The  
 34194 manner in which a mutex is released is dependent upon the mutex's type attribute. If there are  
 34195 threads blocked on the mutex object referenced by *mutex* when *pthread\_mutex\_unlock()* is called,  
 34196 resulting in the mutex becoming available, the scheduling policy shall determine which thread  
 34197 shall acquire the mutex.

34198 XSI (In the case of PTHREAD\_MUTEX\_RECURSIVE mutexes, the mutex shall become available  
 34199 when the count reaches zero and the calling thread no longer has any locks on this mutex).

34200 If a signal is delivered to a thread waiting for a mutex, upon return from the signal handler the  
 34201 thread shall resume waiting for the mutex as if it was not interrupted.

34202 **RETURN VALUE**

34203 If successful, the *pthread\_mutex\_lock()* and *pthread\_mutex\_unlock()* functions shall return zero;  
34204 otherwise, an error number shall be returned to indicate the error.

34205 The *pthread\_mutex\_trylock()* function shall return zero if a lock on the mutex object referenced by  
34206 *mutex* is acquired. Otherwise, an error number is returned to indicate the error.

34207 **ERRORS**

34208 The *pthread\_mutex\_lock()* and *pthread\_mutex\_trylock()* functions shall fail if:

34209 [EINVAL] The *mutex* was created with the protocol attribute having the value  
34210 PTHREAD\_PRIO\_PROTECT and the calling thread's priority is higher than  
34211 the mutex's current priority ceiling.

34212 The *pthread\_mutex\_trylock()* function shall fail if:

34213 [EBUSY] The *mutex* could not be acquired because it was already locked.

34214 The *pthread\_mutex\_lock()*, *pthread\_mutex\_trylock()*, and *pthread\_mutex\_unlock()* functions may  
34215 fail if:

34216 [EINVAL] The value specified by *mutex* does not refer to an initialized mutex object.

34217 XSI [EAGAIN] The mutex could not be acquired because the maximum number of recursive  
34218 locks for *mutex* has been exceeded.

34219 The *pthread\_mutex\_lock()* function may fail if:

34220 [EDEADLK] The current thread already owns the mutex.

34221 The *pthread\_mutex\_unlock()* function may fail if:

34222 [EPERM] The current thread does not own the mutex.

34223 These functions shall not return an error code of [EINTR].

34224 **EXAMPLES**

34225 None.

34226 **APPLICATION USAGE**

34227 None.

34228 **RATIONALE**

34229 Mutex objects are intended to serve as a low-level primitive from which other thread  
34230 synchronization functions can be built. As such, the implementation of mutexes should be as  
34231 efficient as possible, and this has ramifications on the features available at the interface.

34232 The mutex functions and the particular default settings of the mutex attributes have been  
34233 motivated by the desire to not preclude fast, inlined implementations of mutex locking and  
34234 unlocking.

34235 For example, deadlocking on a double-lock is explicitly allowed behavior in order to avoid  
34236 requiring more overhead in the basic mechanism than is absolutely necessary. (More "friendly"  
34237 mutexes that detect deadlock or that allow multiple locking by the same thread are easily  
34238 constructed by the user via the other mechanisms provided. For example, *pthread\_self()* can be  
34239 used to record mutex ownership.) Implementations might also choose to provide such extended  
34240 features as options via special mutex attributes.

34241 Since most attributes only need to be checked when a thread is going to be blocked, the use of  
34242 attributes does not slow the (common) mutex-locking case.

34243 Likewise, while being able to extract the thread ID of the owner of a mutex might be desirable, it  
34244 would require storing the current thread ID when each mutex is locked, and this could incur  
34245 unacceptable levels of overhead. Similar arguments apply to a *mutex\_tryunlock* operation.

34246 **FUTURE DIRECTIONS**

34247 None.

34248 **SEE ALSO**

34249 *pthread\_mutex\_destroy()*, *pthread\_mutex\_timedlock()*, the Base Definitions volume of  
34250 IEEE Std 1003.1-200x, <pthread.h>

34251 **CHANGE HISTORY**

34252 First released in Issue 5. Included for alignment with the POSIX Threads Extension.

34253 **Issue 6**

34254 The *pthread\_mutex\_lock()*, *pthread\_mutex\_trylock()*, and *pthread\_mutex\_unlock()* functions are  
34255 marked as part of the Threads option.

34256 The following new requirements on POSIX implementations derive from alignment with the  
34257 Single UNIX Specification:

- 34258 • The behavior when attempting to relock a mutex is defined.

34259 The *pthread\_mutex\_timedlock()* function is added to the SEE ALSO section for alignment with  
34260 IEEE Std 1003.1d-1999.

34261 **NAME**

34262 pthread\_mutex\_setprioceiling — change the priority ceiling of a mutex (**REALTIME**  
34263 **THREADS**)

34264 **SYNOPSIS**

```
34265 THR TPP #include <pthread.h>
```

```
34266 int pthread_mutex_setprioceiling(pthread_mutex_t *restrict mutex,  
34267 int prioceiling, int *restrict old_ceiling);
```

34268

34269 **DESCRIPTION**

34270 Refer to *pthread\_mutex\_getprioceiling()*.



34271 **NAME**34272 pthread\_mutex\_timedlock — lock a mutex (**ADVANCED REALTIME**)34273 **SYNOPSIS**

34274 THR TMO #include &lt;pthread.h&gt;

34275 #include &lt;time.h&gt;

34276 int pthread\_mutex\_timedlock(pthread\_mutex\_t \*restrict mutex,

34277 const struct timespec \*restrict abs\_timeout);

34278

34279 **DESCRIPTION**

34280 The *pthread\_mutex\_timedlock()* function shall lock the mutex object referenced by *mutex*. If the  
 34281 mutex is already locked, the calling thread shall block until the mutex becomes available as in  
 34282 the *pthread\_mutex\_lock()* function. If the mutex cannot be locked without waiting for another  
 34283 thread to unlock the mutex, this wait shall be terminated when the specified timeout expires.

34284 The timeout shall expire when the absolute time specified by *abs\_timeout* passes, as measured by  
 34285 the clock on which timeouts are based (that is, when the value of that clock equals or exceeds  
 34286 *abs\_timeout*), or if the absolute time specified by *abs\_timeout* has already been passed at the time  
 34287 of the call.

34288 TMR If the Timers option is supported, the timeout shall be based on the `CLOCK_REALTIME` clock; if  
 34289 the Timers option is not supported, the timeout shall be based on the system clock as returned  
 34290 by the *time()* function.

34291 The resolution of the timeout shall be the resolution of the clock on which it is based. The  
 34292 `timespec` data type is defined in the `<time.h>` header.

34293 Under no circumstance shall the function fail with a timeout if the mutex can be locked  
 34294 immediately. The validity of the *abs\_timeout* parameter need not be checked if the mutex can be  
 34295 locked immediately.

34296 As a consequence of the priority inheritance rules (for mutexes initialized with the  
 34297 `PRIO_INHERIT` protocol), if a timed mutex wait is terminated because its timeout expires, the  
 34298 priority of the owner of the mutex shall be adjusted as necessary to reflect the fact that this  
 34299 thread is no longer among the threads waiting for the mutex.

34300 **RETURN VALUE**

34301 If successful, the *pthread\_mutex\_timedlock()* function shall return zero; otherwise, an error  
 34302 number shall be returned to indicate the error.

34303 **ERRORS**

34304 The *pthread\_mutex\_timedlock()* function shall fail if:

34305 [EINVAL] The mutex was created with the protocol attribute having the value  
 34306 `PTHREAD_PRIO_PROTECT` and the calling thread's priority is higher than  
 34307 the mutex' current priority ceiling.

34308 [EINVAL] The process or thread would have blocked, and the *abs\_timeout* parameter  
 34309 specified a nanoseconds field value less than zero or greater than or equal to  
 34310 1 000 million.

34311 [ETIMEDOUT] The mutex could not be locked before the specified timeout expired.

34312 The *pthread\_mutex\_timedlock()* function may fail if:

34313 [EINVAL] The value specified by *mutex* does not refer to an initialized mutex object.

34314 XSI [EAGAIN] The mutex could not be acquired because the maximum number of recursive  
34315 locks for mutex has been exceeded.

34316 [EDEADLK] The current thread already owns the mutex.

34317 This function shall not return an error code of [EINTR].

## 34318 EXAMPLES

34319 None.

## 34320 APPLICATION USAGE

34321 The *pthread\_mutex\_timedlock()* function is part of the Threads and Timeouts options and need  
34322 not be provided on all implementations.

## 34323 RATIONALE

34324 None.

## 34325 FUTURE DIRECTIONS

34326 None.

## 34327 SEE ALSO

34328 *pthread\_mutex\_destroy()*, *pthread\_mutex\_lock()*, *pthread\_mutex\_trylock()*, *time()*, the Base  
34329 Definitions volume of IEEE Std 1003.1-200x, <pthread.h>, <time.h>

## 34330 CHANGE HISTORY

34331 First released in Issue 6. Derived from IEEE Std 1003.1d-1999.

34332 **NAME**

34333 pthread\_mutex\_trylock, pthread\_mutex\_unlock — lock and unlock a mutex

34334 **SYNOPSIS**

34335 THR #include &lt;pthread.h&gt;

34336 int pthread\_mutex\_trylock(pthread\_mutex\_t \*mutex);

34337 int pthread\_mutex\_unlock(pthread\_mutex\_t \*mutex);

34338

34339 **DESCRIPTION**34340 Refer to *pthread\_mutex\_lock()*.

## 34341 NAME

34342 pthread\_mutexattr\_destroy, pthread\_mutexattr\_init — destroy and initialize mutex attributes  
34343 object

## 34344 SYNOPSIS

```
34345 THR #include <pthread.h>
```

```
34346 int pthread_mutexattr_destroy(pthread_mutexattr_t *attr);
```

```
34347 int pthread_mutexattr_init(pthread_mutexattr_t *attr);
```

34348

## 34349 DESCRIPTION

34350 The *pthread\_mutexattr\_destroy()* function shall destroy a mutex attributes object; the object  
34351 becomes, in effect, uninitialized. An implementation may cause *pthread\_mutexattr\_destroy()* to  
34352 set the object referenced by *attr* to an invalid value. A destroyed *attr* attributes object can be  
34353 reinitialized using *pthread\_mutexattr\_init()*; the results of otherwise referencing the object after it  
34354 has been destroyed are undefined. |

34355 The *pthread\_mutexattr\_init()* function shall initialize a mutex attributes object *attr* with the  
34356 default value for all of the attributes defined by the implementation.

34357 Results are undefined if *pthread\_mutexattr\_init()* is called specifying an already initialized *attr* |  
34358 attributes object. |

34359 After a mutex attributes object has been used to initialize one or more mutexes, any function  
34360 affecting the attributes object (including destruction) shall not affect any previously initialized |  
34361 mutexes. |

## 34362 RETURN VALUE

34363 Upon successful completion, *pthread\_mutexattr\_destroy()* and *pthread\_mutexattr\_init()* shall  
34364 return zero; otherwise, an error number shall be returned to indicate the error.

## 34365 ERRORS

34366 The *pthread\_mutexattr\_destroy()* function may fail if:

34367 [EINVAL] The value specified by *attr* is invalid.

34368 The *pthread\_mutexattr\_init()* function shall fail if:

34369 [ENOMEM] Insufficient memory exists to initialize the mutex attributes object.

34370 These functions shall not return an error code of [EINTR].

## 34371 EXAMPLES

34372 None.

## 34373 APPLICATION USAGE

34374 None.

## 34375 RATIONALE

34376 See *pthread\_attr\_init()* for a general explanation of attributes. Attributes objects allow  
34377 implementations to experiment with useful extensions and permit extension of this volume of  
34378 IEEE Std 1003.1-200x without changing the existing functions. Thus, they provide for future  
34379 extensibility of this volume of IEEE Std 1003.1-200x and reduce the temptation to standardize  
34380 prematurely on semantics that are not yet widely implemented or understood.

34381 Examples of possible additional mutex attributes that have been discussed are *spin\_only*,  
34382 *limited\_spin*, *no\_spin*, *recursive*, and *metered*. (To explain what the latter attributes might mean:  
34383 recursive mutexes would allow for multiple re-locking by the current owner; metered mutexes  
34384 would transparently keep records of queue length, wait time, and so on.) Since there is not yet

34385 wide agreement on the usefulness of these resulting from shared implementation and usage  
 34386 experience, they are not yet specified in this volume of IEEE Std 1003.1-200x. Mutex attributes  
 34387 objects, however, make it possible to test out these concepts for possible standardization at a  
 34388 later time.

#### 34389 **Mutex Attributes and Performance**

34390 Care has been taken to ensure that the default values of the mutex attributes have been defined  
 34391 such that mutexes initialized with the defaults have simple enough semantics so that the locking  
 34392 and unlocking can be done with the equivalent of a test-and-set instruction (plus possibly a few  
 34393 other basic instructions).

34394 There is at least one implementation method that can be used to reduce the cost of testing at  
 34395 lock-time if a mutex has non-default attributes. One such method that an implementation can  
 34396 employ (and this can be made fully transparent to fully conforming POSIX applications) is to  
 34397 secretly pre-lock any mutexes that are initialized to non-default attributes. Any later attempt to  
 34398 lock such a mutex causes the implementation to branch to the “slow path” as if the mutex were  
 34399 unavailable; then, on the slow path, the implementation can do the “real work” to lock a non-  
 34400 default mutex. The underlying unlock operation is more complicated since the implementation  
 34401 never really wants to release the pre-lock on this kind of mutex. This illustrates that, depending  
 34402 on the hardware, there may be certain optimizations that can be used so that whatever mutex  
 34403 attributes are considered “most frequently used” can be processed most efficiently.

#### 34404 **Process Shared Memory and Synchronization**

34405 The existence of memory mapping functions in this volume of IEEE Std 1003.1-200x leads to the  
 34406 possibility that an application may allocate the synchronization objects from this section in  
 34407 memory that is accessed by multiple processes (and therefore, by threads of multiple processes).

34408 In order to permit such usage, while at the same time keeping the usual case (that is, usage  
 34409 within a single process) efficient, a process-shared option has been defined.

34410 If an implementation supports the `_POSIX_THREAD_PROCESS_SHARED` option, then the  
 34411 *process-shared* attribute can be used to indicate that mutexes or condition variables may be  
 34412 accessed by threads of multiple processes.

34413 The default setting of `PTHREAD_PROCESS_PRIVATE` has been chosen for the *process-shared*  
 34414 attribute so that the most efficient forms of these synchronization objects are created by default.

34415 Synchronization variables that are initialized with the `PTHREAD_PROCESS_PRIVATE` *process-*  
 34416 *shared* attribute may only be operated on by threads in the process that initialized them.  
 34417 Synchronization variables that are initialized with the `PTHREAD_PROCESS_SHARED` *process-*  
 34418 *shared* attribute may be operated on by any thread in any process that has access to it. In  
 34419 particular, these processes may exist beyond the lifetime of the initializing process. For example,  
 34420 the following code implements a simple counting semaphore in a mapped file that may be used  
 34421 by many processes.

```
34422 /* sem.h */
34423 struct semaphore {
34424     pthread_mutex_t lock;
34425     pthread_cond_t nonzero;
34426     unsigned count;
34427 };
34428 typedef struct semaphore semaphore_t;
34429 semaphore_t *semaphore_create(char *semaphore_name);
34430 semaphore_t *semaphore_open(char *semaphore_name);
```

```
34431 void semaphore_post(semaphore_t *semap);
34432 void semaphore_wait(semaphore_t *semap);
34433 void semaphore_close(semaphore_t *semap);

34434 /* sem.c */
34435 #include <sys/types.h>
34436 #include <sys/stat.h>
34437 #include <sys/mman.h>
34438 #include <fcntl.h>
34439 #include <pthread.h>
34440 #include "sem.h"

34441 semaphore_t *
34442 semaphore_create(char *semaphore_name)
34443 {
34444     int fd;
34445     semaphore_t *semap;
34446     pthread_mutexattr_t psharedm;
34447     pthread_condattr_t psharedc;

34448     fd = open(semaphore_name, O_RDWR | O_CREAT | O_EXCL, 0666);
34449     if (fd < 0)
34450         return (NULL);
34451     (void) ftruncate(fd, sizeof(semaphore_t));
34452     (void) pthread_mutexattr_init(&psharedm);
34453     (void) pthread_mutexattr_setpshared(&psharedm,
34454         PTHREAD_PROCESS_SHARED);
34455     (void) pthread_condattr_init(&psharedc);
34456     (void) pthread_condattr_setpshared(&psharedc,
34457         PTHREAD_PROCESS_SHARED);
34458     semap = (semaphore_t *) mmap(NULL, sizeof(semaphore_t),
34459         PROT_READ | PROT_WRITE, MAP_SHARED,
34460         fd, 0);
34461     close (fd);
34462     (void) pthread_mutex_init(&semap->lock, &psharedm);
34463     (void) pthread_cond_init(&semap->nonzero, &psharedc);
34464     semap->count = 0;
34465     return (semap);
34466 }

34467 semaphore_t *
34468 semaphore_open(char *semaphore_name)
34469 {
34470     int fd;
34471     semaphore_t *semap;

34472     fd = open(semaphore_name, O_RDWR, 0666);
34473     if (fd < 0)
34474         return (NULL);
34475     semap = (semaphore_t *) mmap(NULL, sizeof(semaphore_t),
34476         PROT_READ | PROT_WRITE, MAP_SHARED,
34477         fd, 0);
34478     close (fd);
34479     return (semap);
34480 }
```

```

34481 void
34482 semaphore_post(semaphore_t *semap)
34483 {
34484     pthread_mutex_lock(&semap->lock);
34485     if (semap->count == 0)
34486         pthread_cond_signal(&semap->nonzero);
34487     semap->count++;
34488     pthread_mutex_unlock(&semap->lock);
34489 }
34490 void
34491 semaphore_wait(semaphore_t *semap)
34492 {
34493     pthread_mutex_lock(&semap->lock);
34494     while (semap->count == 0)
34495         pthread_cond_wait(&semap->nonzero, &semap->lock);
34496     semap->count--;
34497     pthread_mutex_unlock(&semap->lock);
34498 }
34499 void
34500 semaphore_close(semaphore_t *semap)
34501 {
34502     munmap((void *) semap, sizeof(semaphore_t));
34503 }

```

34504 The following code is for three separate processes that create, post, and wait on a semaphore in  
34505 the file **/tmp/semaphore**. Once the file is created, the post and wait programs increment and  
34506 decrement the counting semaphore (waiting and waking as required) even though they did not  
34507 initialize the semaphore.

```

34508 /* create.c */
34509 #include "pthread.h"
34510 #include "sem.h"
34511 int
34512 main()
34513 {
34514     semaphore_t *semap;
34515     semap = semaphore_create("/tmp/semaphore");
34516     if (semap == NULL)
34517         exit(1);
34518     semaphore_close(semap);
34519     return (0);
34520 }
34521 /* post */
34522 #include "pthread.h"
34523 #include "sem.h"
34524 int
34525 main()
34526 {
34527     semaphore_t *semap;

```

```
34528     semaphore = semaphore_open("/tmp/semaphore");
34529     if (semap == NULL)
34530         exit(1);
34531     semaphore_post(semap);
34532     semaphore_close(semap);
34533     return (0);
34534 }

34535 /* wait */
34536 #include "pthread.h"
34537 #include "sem.h"

34538 int
34539 main()
34540 {
34541     semaphore_t *semap;

34542     semap = semaphore_open("/tmp/semaphore");
34543     if (semap == NULL)
34544         exit(1);
34545     semaphore_wait(semap);
34546     semaphore_close(semap);
34547     return (0);
34548 }
```

**34549 FUTURE DIRECTIONS**

34550 None.

**34551 SEE ALSO**

34552 *pthread\_cond\_destroy()*, *pthread\_create()*, *pthread\_mutex\_destroy()*, *pthread\_mutexattr\_destroy()*, the  
34553 Base Definitions volume of IEEE Std 1003.1-200x, <pthread.h>

**34554 CHANGE HISTORY**

34555 First released in Issue 5. Included for alignment with the POSIX Threads Extension.

**34556 Issue 6**

34557 The *pthread\_mutexattr\_destroy()* and *pthread\_mutexattr\_init()* functions are marked as part of the  
34558 Threads option.

34559 IEEE PASC Interpretation 1003.1c #27 is applied, updating the ERRORS section.



34560 **NAME**

34561 pthread\_mutexattr\_getprioceiling, pthread\_mutexattr\_setprioceiling — get and set prioceiling  
 34562 attribute of mutex attributes object (**REALTIME THREADS**)

34563 **SYNOPSIS**

```
34564 THR TPP #include <pthread.h>
34565
34566 int pthread_mutexattr_getprioceiling(
34567     const pthread_mutexattr_t *restrict attr,
34568     int *restrict prioceiling);
34569 int pthread_mutexattr_setprioceiling(pthread_mutexattr_t *attr,
34570     int prioceiling);
```

34571 **DESCRIPTION**

34572 The *pthread\_mutexattr\_getprioceiling()* and *pthread\_mutexattr\_setprioceiling()* functions,  
 34573 respectively, shall get and set the priority ceiling attribute of a mutex attributes object pointed to  
 34574 by *attr* which was previously created by the function *pthread\_mutexattr\_init()*.

34575 The *prioceiling* attribute contains the priority ceiling of initialized mutexes. The values of  
 34576 *prioceiling* are within the maximum range of priorities defined by SCHED\_FIFO.

34577 The *prioceiling* attribute defines the priority ceiling of initialized mutexes, which is the minimum  
 34578 priority level at which the critical section guarded by the mutex is executed. In order to avoid  
 34579 priority inversion, the priority ceiling of the mutex shall be set to a priority higher than or equal  
 34580 to the highest priority of all the threads that may lock that mutex. The values of *prioceiling* are  
 34581 within the maximum range of priorities defined under the SCHED\_FIFO scheduling policy.

34582 **RETURN VALUE**

34583 Upon successful completion, the *pthread\_mutexattr\_getprioceiling()* and  
 34584 *pthread\_mutexattr\_setprioceiling()* functions shall return zero; otherwise, an error number shall be  
 34585 returned to indicate the error.

34586 **ERRORS**

34587 The *pthread\_mutexattr\_getprioceiling()* and *pthread\_mutexattr\_setprioceiling()* functions may fail if:

34588 [EINVAL] The value specified by *attr* or *prioceiling* is invalid.

34589 [EPERM] The caller does not have the privilege to perform the operation.

34590 These functions shall not return an error code of [EINTR].

34591 **EXAMPLES**

34592 None.

34593 **APPLICATION USAGE**

34594 None.

34595 **RATIONALE**

34596 None.

34597 **FUTURE DIRECTIONS**

34598 None.

34599 **SEE ALSO**

34600 *pthread\_cond\_destroy()*, *pthread\_create()*, *pthread\_mutex\_destroy()*, the Base Definitions volume of  
 34601 IEEE Std 1003.1-200x, <pthread.h>

## 34602 CHANGE HISTORY

34603 First released in Issue 5. Included for alignment with the POSIX Threads Extension.

34604 Marked as part of the Realtime Threads Feature Group.

## 34605 Issue 6

34606 The `pthread_mutexattr_getprioceiling()` and `pthread_mutexattr_setprioceiling()` functions are  
34607 marked as part of the Threads and Thread Priority Protection options.

34608 The [ENOSYS] error condition has been removed as stubs need not be provided if an  
34609 implementation does not support the Thread Priority Protection option.

34610 The [ENOTSUP] error condition has been removed since these functions do not have a *protocol*  
34611 argument.

34612 The `restrict` keyword is added to the `pthread_mutexattr_getprioceiling()` prototype for alignment  
34613 with the ISO/IEC 9899:1999 standard.

## 34614 NAME

34615 pthread\_mutexattr\_getprotocol, pthread\_mutexattr\_setprotocol — get and set protocol attribute  
 34616 of mutex attributes object (**REALTIME THREADS**)

## 34617 SYNOPSIS

```
34618 THR      #include <pthread.h>
34619 TPP|TPI   int pthread_mutexattr_getprotocol(
34620           const pthread_mutexattr_t *restrict attr,
34621           int *restrict protocol);
34622           int pthread_mutexattr_setprotocol(pthread_mutexattr_t *attr,
34623           int protocol);
34624
```

## 34625 DESCRIPTION

34626 The *pthread\_mutexattr\_getprotocol()* and *pthread\_mutexattr\_setprotocol()* functions, respectively,  
 34627 shall get and set the protocol attribute of a mutex attributes object pointed to by *attr* which was  
 34628 previously created by the function *pthread\_mutexattr\_init()*.

34629 The *protocol* attribute defines the protocol to be followed in utilizing mutexes. The value of  
 34630 *protocol* may be one of:

```
34631          PTHREAD_PRIO_NONE
34632 TPI      PTHREAD_PRIO_INHERIT
34633 TPP      PTHREAD_PRIO_PROTECT
34634
```

34635 which are defined in the **<pthread.h>** header. |

34636 When a thread owns a mutex with the PTHREAD\_PRIO\_NONE *protocol* attribute, its priority |  
 34637 and scheduling shall not be affected by its mutex ownership. |

34638 TPI When a thread is blocking higher priority threads because of owning one or more mutexes with  
 34639 the PTHREAD\_PRIO\_INHERIT protocol attribute, it shall execute at the higher of its priority or |  
 34640 the priority of the highest priority thread waiting on any of the mutexes owned by this thread |  
 34641 and initialized with this protocol. |

34642 TPP When a thread owns one or more mutexes initialized with the PTHREAD\_PRIO\_PROTECT |  
 34643 protocol, it shall execute at the higher of its priority or the highest of the priority ceilings of all |  
 34644 the mutexes owned by this thread and initialized with this attribute, regardless of whether other |  
 34645 threads are blocked on any of these mutexes or not. |

34646 While a thread is holding a mutex which has been initialized with the  
 34647 PTHREAD\_PRIO\_INHERIT or PTHREAD\_PRIO\_PROTECT protocol attributes, it shall not be  
 34648 subject to being moved to the tail of the scheduling queue at its priority in the event that its  
 34649 original priority is changed, such as by a call to *sched\_setparam()*. Likewise, when a thread  
 34650 unlocks a mutex that has been initialized with the PTHREAD\_PRIO\_INHERIT or  
 34651 PTHREAD\_PRIO\_PROTECT protocol attributes, it shall not be subject to being moved to the tail  
 34652 of the scheduling queue at its priority in the event that its original priority is changed.

34653 If a thread simultaneously owns several mutexes initialized with different protocols, it shall  
 34654 execute at the highest of the priorities that it would have obtained by each of these protocols.

34655 TPI When a thread makes a call to *pthread\_mutex\_lock()*, the mutex was initialized with the protocol  
 34656 attribute having the value PTHREAD\_PRIO\_INHERIT, when the calling thread is blocked  
 34657 because the mutex is owned by another thread, that owner thread shall inherit the priority level  
 34658 of the calling thread as long as it continues to own the mutex. The implementation shall update  
 34659 its execution priority to the maximum of its assigned priority and all its inherited priorities.

34660 Furthermore, if this owner thread itself becomes blocked on another mutex, the same priority  
34661 inheritance effect shall be propagated to this other owner thread, in a recursive manner.

## 34662 RETURN VALUE

34663 Upon successful completion, the *pthread\_mutexattr\_getprotocol()* and  
34664 *pthread\_mutexattr\_setprotocol()* functions shall return zero; otherwise, an error number shall be  
34665 returned to indicate the error.

## 34666 ERRORS

34667 The *pthread\_mutexattr\_setprotocol()* function shall fail if:

34668 [ENOTSUP] The value specified by *protocol* is an unsupported value.

34669 The *pthread\_mutexattr\_getprotocol()* and *pthread\_mutexattr\_setprotocol()* functions may fail if:

34670 [EINVAL] The value specified by *attr* or *protocol* is invalid.

34671 [EPERM] The caller does not have the privilege to perform the operation.

34672 These functions shall not return an error code of [EINTR].

## 34673 EXAMPLES

34674 None.

## 34675 APPLICATION USAGE

34676 None.

## 34677 RATIONALE

34678 None.

## 34679 FUTURE DIRECTIONS

34680 None.

## 34681 SEE ALSO

34682 *pthread\_cond\_destroy()*, *pthread\_create()*, *pthread\_mutex\_destroy()*, the Base Definitions volume of  
34683 IEEE Std 1003.1-200x, <pthread.h>

## 34684 CHANGE HISTORY

34685 First released in Issue 5. Included for alignment with the POSIX Threads Extension.

34686 Marked as part of the Realtime Threads Feature Group.

## 34687 Issue 6

34688 The *pthread\_mutexattr\_getprotocol()* and *pthread\_mutexattr\_setprotocol()* functions are marked as  
34689 part of the Threads option and either the Thread Priority Protection or Thread Priority  
34690 Inheritance options.

34691 The [ENOSYS] error condition has been removed as stubs need not be provided if an  
34692 implementation does not support the Thread Priority Protection or Thread Priority Inheritance  
34693 options.

34694 The **restrict** keyword is added to the *pthread\_mutexattr\_getprotocol()* prototype for alignment  
34695 with the ISO/IEC 9899:1999 standard.

34696 **NAME**

34697 pthread\_mutexattr\_getpshared, pthread\_mutexattr\_setpshared — get and set process-shared  
 34698 attribute

34699 **SYNOPSIS**

```
34700 THR TSH #include <pthread.h>
```

```
34701 int pthread_mutexattr_getpshared(  
34702     const pthread_mutexattr_t *restrict attr,  
34703     int *restrict pshared);  
34704 int pthread_mutexattr_setpshared(pthread_mutexattr_t *attr,  
34705     int pshared);  
34706
```

34707 **DESCRIPTION**

34708 The *pthread\_mutexattr\_getpshared()* function shall obtain the value of the *process-shared* attribute  
 34709 from the attributes object referenced by *attr*. The *pthread\_mutexattr\_setpshared()* function shall  
 34710 set the *process-shared* attribute in an initialized attributes object referenced by *attr*.

34711 The *process-shared* attribute is set to PTHREAD\_PROCESS\_SHARED to permit a mutex to be  
 34712 operated upon by any thread that has access to the memory where the mutex is allocated, even if  
 34713 the mutex is allocated in memory that is shared by multiple processes. If the *process-shared*  
 34714 attribute is PTHREAD\_PROCESS\_PRIVATE, the mutex shall only be operated upon by threads  
 34715 created within the same process as the thread that initialized the mutex; if threads of differing  
 34716 processes attempt to operate on such a mutex, the behavior is undefined. The default value of  
 34717 the attribute shall be PTHREAD\_PROCESS\_PRIVATE.

34718 **RETURN VALUE**

34719 Upon successful completion, *pthread\_mutexattr\_setpshared()* shall return zero; otherwise, an error  
 34720 number shall be returned to indicate the error.

34721 Upon successful completion, *pthread\_mutexattr\_getpshared()* shall return zero and stores the  
 34722 value of the *process-shared* attribute of *attr* into the object referenced by the *pshared* parameter.  
 34723 Otherwise, an error number shall be returned to indicate the error.

34724 **ERRORS**

34725 The *pthread\_mutexattr\_getpshared()* and *pthread\_mutexattr\_setpshared()* functions may fail if:

34726 [EINVAL] The value specified by *attr* is invalid.

34727 The *pthread\_mutexattr\_setpshared()* function may fail if:

34728 [EINVAL] The new value specified for the attribute is outside the range of legal values  
 34729 for that attribute.

34730 These functions shall not return an error code of [EINTR].

34731 **EXAMPLES**

34732 None.

34733 **APPLICATION USAGE**

34734 None.

34735 **RATIONALE**

34736 None.

34737 **FUTURE DIRECTIONS**

34738 None.

34739 **SEE ALSO**

34740 *pthread\_cond\_destroy()*, *pthread\_create()*, *pthread\_mutex\_destroy()*, *pthread\_mutexattr\_destroy()*, the  
34741 Base Definitions volume of IEEE Std 1003.1-200x, <**pthread.h**>

34742 **CHANGE HISTORY**

34743 First released in Issue 5. Included for alignment with the POSIX Threads Extension.

34744 **Issue 6**

34745 The *pthread\_mutexattr\_getpshared()* and *pthread\_mutexattr\_setpshared()* functions are marked as  
34746 part of the Threads and Thread Process-Shared Synchronization options.

34747 The **restrict** keyword is added to the *pthread\_mutexattr\_getpshared()* prototype for alignment  
34748 with the ISO/IEC 9899:1999 standard.

34749 **NAME**

34750 pthread\_mutexattr\_gettype, pthread\_mutexattr\_settype — get and set a mutex type attribute

34751 **SYNOPSIS**

```
34752 XSI #include <pthread.h>
34753
34754 int pthread_mutexattr_gettype(const pthread_mutexattr_t *restrict attr,
34755 int *restrict type);
34756 int pthread_mutexattr_settype(pthread_mutexattr_t *attr, int type);
```

34757 **DESCRIPTION**

34758 The *pthread\_mutexattr\_gettype()* and *pthread\_mutexattr\_settype()* functions, respectively, shall get  
 34759 and set the mutex *type* attribute. This attribute is set in the *type* parameter to these functions. The  
 34760 default value of the *type* attribute is PTHREAD\_MUTEX\_DEFAULT.

34761 The type of mutex is contained in the *type* attribute of the mutex attributes. Valid mutex types  
 34762 include:

## 34763 PTHREAD\_MUTEX\_NORMAL

34764 This type of mutex does not detect deadlock. A thread attempting to relock this mutex  
 34765 without first unlocking it shall deadlock. Attempting to unlock a mutex locked by a  
 34766 different thread results in undefined behavior. Attempting to unlock an unlocked mutex  
 34767 results in undefined behavior.

## 34768 PTHREAD\_MUTEX\_ERRORCHECK

34769 This type of mutex provides error checking. A thread attempting to relock this mutex  
 34770 without first unlocking it shall return with an error. A thread attempting to unlock a mutex  
 34771 which another thread has locked shall return with an error. A thread attempting to unlock  
 34772 an unlocked mutex shall return with an error.

## 34773 PTHREAD\_MUTEX\_RECURSIVE

34774 A thread attempting to relock this mutex without first unlocking it shall succeed in locking  
 34775 the mutex. The relocking deadlock which can occur with mutexes of type  
 34776 PTHREAD\_MUTEX\_NORMAL cannot occur with this type of mutex. Multiple locks of this  
 34777 mutex shall require the same number of unlocks to release the mutex before another thread  
 34778 can acquire the mutex. A thread attempting to unlock a mutex which another thread has  
 34779 locked shall return with an error. A thread attempting to unlock an unlocked mutex shall  
 34780 return with an error.

## 34781 PTHREAD\_MUTEX\_DEFAULT

34782 Attempting to recursively lock a mutex of this type results in undefined behavior.  
 34783 Attempting to unlock a mutex of this type which was not locked by the calling thread  
 34784 results in undefined behavior. Attempting to unlock a mutex of this type which is not  
 34785 locked results in undefined behavior. An implementation may map this mutex to one of the  
 34786 other mutex types.

34787 **RETURN VALUE**

34788 Upon successful completion, the *pthread\_mutexattr\_gettype()* function shall return zero and store  
 34789 the value of the *type* attribute of *attr* into the object referenced by the *type* parameter. Otherwise,  
 34790 an error shall be returned to indicate the error.

34791 If successful, the *pthread\_mutexattr\_settype()* function shall return zero; otherwise, an error  
 34792 number shall be returned to indicate the error.

34793 **ERRORS**

34794 The *pthread\_mutexattr\_settype()* function shall fail if:

34795 [EINVAL] The value *type* is invalid.

34796 The *pthread\_mutexattr\_gettype()* and *pthread\_mutexattr\_settype()* functions may fail if:

34797 [EINVAL] The value specified by *attr* is invalid.

34798 These functions shall not return an error code of [EINTR].

34799 **EXAMPLES**

34800 None.

34801 **APPLICATION USAGE**

34802 It is advised that an application should not use a PTHREAD\_MUTEX\_RECURSIVE mutex with  
34803 condition variables because the implicit unlock performed for a *pthread\_cond\_timedwait()* or  
34804 *pthread\_cond\_wait()* may not actually release the mutex (if it had been locked multiple times). If  
34805 this happens, no other thread can satisfy the condition of the predicate.

34806 **RATIONALE**

34807 None.

34808 **FUTURE DIRECTIONS**

34809 None.

34810 **SEE ALSO**

34811 *pthread\_cond\_timedwait()*, *pthread\_cond\_wait()*, the Base Definitions volume of  
34812 IEEE Std 1003.1-200x, <pthread.h>

34813 **CHANGE HISTORY**

34814 First released in Issue 5.

34815 **Issue 6**

34816 The Open Group Corrigendum U033/3 is applied. The SYNOPSIS for  
34817 *pthread\_mutexattr\_gettype()* is updated so that the first argument is of type **const**  
34818 **pthread\_mutexattr\_t** \*.

34819 The **restrict** keyword is added to the *pthread\_mutexattr\_gettype()* prototype for alignment with  
34820 the ISO/IEC 9899:1999 standard.



34821 **NAME**

34822 pthread\_mutexattr\_init — initialize mutex attributes object

34823 **SYNOPSIS**

34824 THR #include &lt;pthread.h&gt;

34825 int pthread\_mutexattr\_init(pthread\_mutexattr\_t \*attr);

34826

34827 **DESCRIPTION**34828 Refer to *pthread\_mutexattr\_destroy()*.

34829 **NAME**

34830 pthread\_mutexattr\_setprioceiling — set prioceiling attribute of mutex attributes object  
34831 (**REALTIME THREADS**)

34832 **SYNOPSIS**

34833 THR TPP #include <pthread.h>

```
34834 int pthread_mutexattr_setprioceiling(pthread_mutexattr_t *attr,  
34835 int prioceiling);
```

34836

34837 **DESCRIPTION**

34838 Refer to *pthread\_mutexattr\_getprioceiling()*.

34839 **NAME**

34840 pthread\_mutexattr\_setprotocol — set protocol attribute of mutex attributes object (**REALTIME**  
34841 **THREADS**)

34842 **SYNOPSIS**

34843 THR #include <pthread.h>

34844 TPP|TPI int pthread\_mutexattr\_setprotocol(pthread\_mutexattr\_t \*attr,  
34845 int protocol);

34846

34847 **DESCRIPTION**

34848 Refer to *pthread\_mutexattr\_setprotocol()*.

34849 **NAME**

34850 pthread\_mutexattr\_setpshared — set process-shared attribute

34851 **SYNOPSIS**

34852 THR TSH #include <pthread.h>

```
34853 int pthread_mutexattr_setpshared(pthread_mutexattr_t *attr,  
34854 int pshared);
```

34855

34856 **DESCRIPTION**

34857 Refer to *pthread\_mutexattr\_getpshared()*.

34858 **NAME**

34859 pthread\_mutexattr\_settype — set a mutex type attribute

34860 **SYNOPSIS**34861 XSI `#include <pthread.h>`34862 `int pthread_mutexattr_settype(pthread_mutexattr_t *attr, int type);`

34863

34864 **DESCRIPTION**34865 Refer to *pthread\_mutexattr\_gettype()*.

## 34866 NAME

34867 pthread\_once — dynamic package initialization

## 34868 SYNOPSIS

34869 THR #include &lt;pthread.h&gt;

```

34870 int pthread_once(pthread_once_t *once_control,
34871                 void (*init_routine)(void));
34872 pthread_once_t once_control = PTHREAD_ONCE_INIT;
34873 
```

## 34874 DESCRIPTION

34875 The first call to *pthread\_once()* by any thread in a process, with a given *once\_control*, shall call the  
 34876 *init\_routine* with no arguments. Subsequent calls of *pthread\_once()* with the same *once\_control*  
 34877 shall not call the *init\_routine*. On return from *pthread\_once()*, *init\_routine* shall have completed. |  
 34878 The *once\_control* parameter shall determine whether the associated initialization routine has |  
 34879 been called.

34880 The *pthread\_once()* function is not a cancellation point. However, if *init\_routine* is a cancellation  
 34881 point and is canceled, the effect on *once\_control* shall be as if *pthread\_once()* was never called.

34882 The constant PTHREAD\_ONCE\_INIT is defined in the <pthread.h> header. |

34883 The behavior of *pthread\_once()* is undefined if *once\_control* has automatic storage duration or is  
 34884 not initialized by PTHREAD\_ONCE\_INIT.

## 34885 RETURN VALUE

34886 Upon successful completion, *pthread\_once()* shall return zero; otherwise, an error number shall  
 34887 be returned to indicate the error.

## 34888 ERRORS

34889 The *pthread\_once()* function may fail if:

34890 [EINVAL] If either *once\_control* or *init\_routine* is invalid.

34891 The *pthread\_once()* function shall not return an error code of [EINTR].

## 34892 EXAMPLES

34893 None.

## 34894 APPLICATION USAGE

34895 None.

## 34896 RATIONALE

34897 Some C libraries are designed for dynamic initialization. That is, the global initialization for the  
 34898 library is performed when the first procedure in the library is called. In a single-threaded  
 34899 program, this is normally implemented using a static variable whose value is checked on entry  
 34900 to a routine, as follows:

```

34901 static int random_is_initialized = 0;
34902 extern int initialize_random();

34903 int random_function()
34904 {
34905     if (random_is_initialized == 0) {
34906         initialize_random();
34907         random_is_initialized = 1;
34908     }
34909     ... /* Operations performed after initialization. */
34910 }
```

34911 To keep the same structure in a multi-threaded program, a new primitive is needed. Otherwise,  
34912 library initialization has to be accomplished by an explicit call to a library-exported initialization  
34913 function prior to any use of the library.

34914 For dynamic library initialization in a multi-threaded process, a simple initialization flag is not  
34915 sufficient; the flag needs to be protected against modification by multiple threads  
34916 simultaneously calling into the library. Protecting the flag requires the use of a mutex; however,  
34917 mutexes have to be initialized before they are used. Ensuring that the mutex is only initialized  
34918 once requires a recursive solution to this problem.

34919 The use of *pthread\_once()* not only supplies an implementation-guaranteed means of dynamic  
34920 initialization, it provides an aid to the reliable construction of multi-threaded and realtime  
34921 systems. The preceding example then becomes:

```
34922 #include <pthread.h>
34923 static pthread_once_t random_is_initialized = PTHREAD_ONCE_INIT;
34924 extern int initialize_random();

34925 int random_function()
34926 {
34927     (void) pthread_once(&random_is_initialized, initialize_random);
34928     ... /* Operations performed after initialization. */
34929 }
```

34930 Note that a **pthread\_once\_t** cannot be an array because some compilers do not accept the  
34931 construct **&<array\_name>**.

#### 34932 **FUTURE DIRECTIONS**

34933 None.

#### 34934 **SEE ALSO**

34935 The Base Definitions volume of IEEE Std 1003.1-200x, **<pthread.h>**

#### 34936 **CHANGE HISTORY**

34937 First released in Issue 5. Included for alignment with the POSIX Threads Extension.

#### 34938 **Issue 6**

34939 The *pthread\_once()* function is marked as part of the Threads option.

34940 The [EINVAL] error is added as a may fail case for if either argument is invalid.

## 34941 NAME

34942 pthread\_rwlock\_destroy, pthread\_rwlock\_init — destroy and initialize a read-write lock object

## 34943 SYNOPSIS

34944 THR #include &lt;pthread.h&gt;

```
34945 int pthread_rwlock_destroy(pthread_rwlock_t *rwlock);
34946 int pthread_rwlock_init(pthread_rwlock_t *restrict rwlock,
34947 const pthread_rwlockattr_t *restrict attr);
34948
```

## 34949 DESCRIPTION

34950 The *pthread\_rwlock\_destroy()* function shall destroy the read-write lock object referenced by  
 34951 *rwlock* and release any resources used by the lock. The effect of subsequent use of the lock is  
 34952 undefined until the lock is reinitialized by another call to *pthread\_rwlock\_init()*. An  
 34953 implementation may cause *pthread\_rwlock\_destroy()* to set the object referenced by *rwlock* to an  
 34954 invalid value. Results are undefined if *pthread\_rwlock\_destroy()* is called when any thread holds  
 34955 *rwlock*. Attempting to destroy an uninitialized read-write lock results in undefined behavior.

34956 The *pthread\_rwlock\_init()* function shall allocate any resources required to use the read-write  
 34957 lock referenced by *rwlock* and initializes the lock to an unlocked state with attributes referenced  
 34958 by *attr*. If *attr* is NULL, the default read-write lock attributes shall be used; the effect is the same  
 34959 as passing the address of a default read-write lock attributes object. Once initialized, the lock can  
 34960 be used any number of times without being reinitialized. Results are undefined if  
 34961 *pthread\_rwlock\_init()* is called specifying an already initialized read-write lock. Results are  
 34962 undefined if a read-write lock is used without first being initialized.

34963 If the *pthread\_rwlock\_init()* function fails, *rwlock* shall not be initialized and the contents of *rwlock*  
 34964 are undefined.

34965 Only the object referenced by *rwlock* may be used for performing synchronization. The result of  
 34966 referring to copies of that object in calls to *pthread\_rwlock\_destroy()*, *pthread\_rwlock\_rdlock()*,  
 34967 *pthread\_rwlock\_timedrdlock()*, *pthread\_rwlock\_timedwrlock()*, *pthread\_rwlock\_tryrdlock()*,  
 34968 *pthread\_rwlock\_trywrlock()*, *pthread\_rwlock\_unlock()*, or *pthread\_rwlock\_wrlock()* is undefined.

## 34969 RETURN VALUE

34970 If successful, the *pthread\_rwlock\_destroy()* and *pthread\_rwlock\_init()* functions shall return zero;  
 34971 otherwise, an error number shall be returned to indicate the error.

34972 The [EBUSY] and [EINVAL] error checks, if implemented, act as if they were performed  
 34973 immediately at the beginning of processing for the function and caused an error return prior to  
 34974 modifying the state of the read-write lock specified by *rwlock*.

## 34975 ERRORS

34976 The *pthread\_rwlock\_destroy()* function may fail if:

34977 [EBUSY] The implementation has detected an attempt to destroy the object referenced  
 34978 by *rwlock* while it is locked.

34979 [EINVAL] The value specified by *rwlock* is invalid.

34980 The *pthread\_rwlock\_init()* function shall fail if:

34981 [EAGAIN] The system lacked the necessary resources (other than memory) to initialize  
 34982 another read-write lock.

34983 [ENOMEM] Insufficient memory exists to initialize the read-write lock.

34984 [EPERM] The caller does not have the privilege to perform the operation.



- 34985 The *pthread\_rwlock\_init()* function may fail if:
- 34986 [EBUSY] The implementation has detected an attempt to reinitialize the object  
34987 referenced by *rwlock*, a previously initialized but not yet destroyed read-write  
34988 lock.
- 34989 [EINVAL] The value specified by *attr* is invalid.
- 34990 These functions shall not return an error code of [EINTR].
- 34991 **EXAMPLES**
- 34992 None.
- 34993 **APPLICATION USAGE**
- 34994 None.
- 34995 **RATIONALE**
- 34996 None.
- 34997 **FUTURE DIRECTIONS**
- 34998 None.
- 34999 **SEE ALSO**
- 35000 *pthread\_rwlock\_rdlock()*, *pthread\_rwlock\_timedrdlock()*, *pthread\_rwlock\_timedwrlock()*,  
35001 *pthread\_rwlock\_tryrdlock()*, *pthread\_rwlock\_trywrlock()*, *pthread\_rwlock\_unlock()*,  
35002 *pthread\_rwlock\_wrlock()*, the Base Definitions volume of IEEE Std 1003.1-200x, <pthread.h>
- 35003 **CHANGE HISTORY**
- 35004 First released in Issue 5.
- 35005 **Issue 6**
- 35006 The following changes are made for alignment with IEEE Std 1003.1j-2000:
- 35007 • The margin code in the SYNOPSIS is changed to THR to indicate that the functionality is  
35008 now part of the Threads option (previously it was part of the Read-Write Locks option in  
35009 IEEE Std 1003.1j-2000 and also part of the XSI extension). The initializer macro is also deleted  
35010 from the SYNOPSIS.
- 35011 • The DESCRIPTION is updated as follows:
- 35012 — It explicitly notes allocation of resources upon initialization of a read-write lock object.
- 35013 — A paragraph is added specifying that copies of read-write lock objects may not be used.
- 35014 • An [EINVAL] error is added to the ERRORS section for *pthread\_rwlock\_init()*, indicating that  
35015 the *rwlock* value is invalid.
- 35016 • The SEE ALSO section is updated.
- 35017 The **restrict** keyword is added to the *pthread\_rwlock\_init()* prototype for alignment with the  
35018 ISO/IEC 9899:1999 standard.

35019 **NAME**

35020 pthread\_rwlock\_init — initialize a read-write lock object

35021 **SYNOPSIS**

35022 THR #include <pthread.h>

```
35023 int pthread_rwlock_init(pthread_rwlock_t *restrict rwlock,  
35024 const pthread_rwlockattr_t *restrict attr);  
35025
```

35026 **DESCRIPTION**

35027 Refer to *pthread\_rwlock\_destroy()*.

35028 **NAME**

35029 pthread\_rwlock\_rdlock, pthread\_rwlock\_tryrdlock — lock a read-write lock object for reading

35030 **SYNOPSIS**

35031 THR #include &lt;pthread.h&gt;

35032 int pthread\_rwlock\_rdlock(pthread\_rwlock\_t \*rwlock);

35033 int pthread\_rwlock\_tryrdlock(pthread\_rwlock\_t \*rwlock);

35034

35035 **DESCRIPTION**

35036 The *pthread\_rwlock\_rdlock()* function shall apply a read lock to the read-write lock referenced by  
 35037 *rwlock*. The calling thread acquires the read lock if a writer does not hold the lock and there are  
 35038 no writers blocked on the lock.

35039 TPS If the Thread Execution Scheduling option is supported, and the threads involved in the lock are  
 35040 executing with the scheduling policies SCHED\_FIFO or SCHED\_RR, the calling thread shall not  
 35041 acquire the lock if a writer holds the lock or if writers of higher or equal priority are blocked on  
 35042 the lock; otherwise, the calling thread shall acquire the lock.

35043 TPS TSP If the Threads Execution Scheduling option is supported, and the threads involved in the lock  
 35044 are executing with the SCHED\_SPORADIC scheduling policy, the calling thread shall not  
 35045 acquire the lock if a writer holds the lock or if writers of higher or equal priority are blocked on  
 35046 the lock; otherwise, the calling thread shall acquire the lock.

35047 If the Thread Execution Scheduling option is not supported, it is implementation-defined  
 35048 whether the calling thread acquires the lock when a writer does not hold the lock and there are  
 35049 writers blocked on the lock. If a writer holds the lock, the calling thread shall not acquire the  
 35050 read lock. If the read lock is not acquired, the calling thread shall block until it can acquire the  
 35051 lock. The calling thread may deadlock if at the time the call is made it holds a write lock.

35052 A thread may hold multiple concurrent read locks on *rwlock* (that is, successfully call the  
 35053 *pthread\_rwlock\_rdlock()* function *n* times). If so, the application shall ensure that the thread  
 35054 performs matching unlocks (that is, it calls the *pthread\_rwlock\_unlock()* function *n* times).

35055 The maximum number of simultaneous read locks that an implementation guarantees can be  
 35056 applied to a read-write lock shall be implementation-defined. The *pthread\_rwlock\_rdlock()*  
 35057 function may fail if this maximum would be exceeded.

35058 The *pthread\_rwlock\_tryrdlock()* function shall apply a read lock as in the *pthread\_rwlock\_rdlock()*  
 35059 function, with the exception that the function shall fail if the equivalent *pthread\_rwlock\_rdlock()*  
 35060 call would have blocked the calling thread. In no case shall the *pthread\_rwlock\_tryrdlock()*  
 35061 function ever block; it always either acquires the lock or fails and returns immediately.

35062 Results are undefined if any of these functions are called with an uninitialized read-write lock.

35063 If a signal is delivered to a thread waiting for a read-write lock for reading, upon return from the  
 35064 signal handler the thread resumes waiting for the read-write lock for reading as if it was not  
 35065 interrupted.

35066 **RETURN VALUE**

35067 If successful, the *pthread\_rwlock\_rdlock()* function shall return zero; otherwise, an error number  
 35068 shall be returned to indicate the error.

35069 The *pthread\_rwlock\_tryrdlock()* function shall return zero if the lock for reading on the read-write  
 35070 lock object referenced by *rwlock* is acquired. Otherwise, an error number shall be returned to  
 35071 indicate the error.

35072 **ERRORS**

35073 The *pthread\_rwlock\_tryrdlock()* function shall fail if:

35074 [EBUSY] The read-write lock could not be acquired for reading because a writer holds  
35075 the lock or a writer with the appropriate priority was blocked on it.

35076 The *pthread\_rwlock\_rdlock()* and *pthread\_rwlock\_tryrdlock()* functions may fail if:

35077 [EINVAL] The value specified by *rwlock* does not refer to an initialized read-write lock  
35078 object.

35079 [EAGAIN] The read lock could not be acquired because the maximum number of read  
35080 locks for *rwlock* has been exceeded.

35081 The *pthread\_rwlock\_rdlock()* function may fail if:

35082 [EDEADLK] The current thread already owns the read-write lock for writing.

35083 These functions shall not return an error code of [EINTR].

35084 **EXAMPLES**

35085 None.

35086 **APPLICATION USAGE**

35087 Applications using these functions may be subject to priority inversion, as discussed in the Base  
35088 Definitions volume of IEEE Std 1003.1-200x, Section 3.285, Priority Inversion.

35089 **RATIONALE**

35090 None.

35091 **FUTURE DIRECTIONS**

35092 None.

35093 **SEE ALSO**

35094 *pthread\_rwlock\_destroy()*, *pthread\_rwlock\_init()*, *pthread\_rwlock\_timedrdlock()*,  
35095 *pthread\_rwlock\_timedwrlock()*, *pthread\_rwlock\_trywrlock()*, *pthread\_rwlock\_unlock()*,  
35096 *pthread\_rwlock\_wrlock()*, the Base Definitions volume of IEEE Std 1003.1-200x, <**pthread.h**>

35097 **CHANGE HISTORY**

35098 First released in Issue 5.

35099 **Issue 6**

35100 The following changes are made for alignment with IEEE Std 1003.1j-2000:

35101 • The margin code in the SYNOPSIS is changed to THR to indicate that the functionality is  
35102 now part of the Threads option (previously it was part of the Read-Write Locks option in  
35103 IEEE Std 1003.1j-2000 and also part of the XSI extension).

35104 • The DESCRIPTION is updated as follows:

35105 — Conditions under which writers have precedence over readers are specified.

35106 — Failure of *pthread\_rwlock\_tryrdlock()* is clarified.

35107 — A paragraph on the maximum number of read locks is added.

35108 • In the ERRORS sections, [EBUSY] is modified to take into account write priority, and  
35109 [EDEADLK] is deleted as a *pthread\_rwlock\_tryrdlock()* error.

35110 • The SEE ALSO section is updated.

35111 **NAME**

35112 pthread\_rwlock\_timedrdlock — lock a read-write lock for reading

35113 **SYNOPSIS**

35114 THR TMO #include &lt;pthread.h&gt;

35115 #include &lt;time.h&gt;

35116 int pthread\_rwlock\_timedrdlock(pthread\_rwlock\_t \*restrict *rwlock*,35117 const struct timespec \*restrict *abs\_timeout*);

35118

35119 **DESCRIPTION**

35120 The *pthread\_rwlock\_timedrdlock()* function shall apply a read lock to the read-write lock  
 35121 referenced by *rwlock* as in the *pthread\_rwlock\_rdlock()* function. However, if the lock cannot be  
 35122 acquired without waiting for other threads to unlock the lock, this wait shall be terminated  
 35123 when the specified timeout expires. The timeout shall expire when the absolute time specified  
 35124 by *abs\_timeout* passes, as measured by the clock on which timeouts are based (that is, when the  
 35125 value of that clock equals or exceeds *abs\_timeout*), or if the absolute time specified by *abs\_timeout*  
 35126 has already been passed at the time of the call.

35127 TMR If the Timers option is supported, the timeout shall be based on the `CLOCK_REALTIME` clock. If  
 35128 the Timers option is not supported, the timeout shall be based on the system clock as returned  
 35129 by the *time()* function. The resolution of the timeout shall be the resolution of the clock on which  
 35130 it is based. The **timespec** data type is defined in the `<time.h>` header. Under no circumstances  
 35131 shall the function fail with a timeout if the lock can be acquired immediately. The validity of the  
 35132 *abs\_timeout* parameter need not be checked if the lock can be immediately acquired.

35133 If a signal that causes a signal handler to be executed is delivered to a thread blocked on a read-  
 35134 write lock via a call to *pthread\_rwlock\_timedrdlock()*, upon return from the signal handler the  
 35135 thread shall resume waiting for the lock as if it was not interrupted.

35136 The calling thread may deadlock if at the time the call is made it holds a write lock on *rwlock*.  
 35137 The results are undefined if this function is called with an uninitialized read-write lock.

35138 **RETURN VALUE**

35139 The *pthread\_rwlock\_timedrdlock()* function shall return zero if the lock for reading on the read-  
 35140 write lock object referenced by *rwlock* is acquired. Otherwise, an error number shall be returned  
 35141 to indicate the error.

35142 **ERRORS**35143 The *pthread\_rwlock\_timedrdlock()* function shall fail if:

35144 [ETIMEDOUT] The lock could not be acquired before the specified timeout expired.

35145 The *pthread\_rwlock\_timedrdlock()* function may fail if:35146 [EAGAIN] The read lock could not be acquired because the maximum number of read  
35147 locks for lock would be exceeded.35148 [EDEADLK] The calling thread already holds a write lock on *rwlock*.35149 [EINVAL] The value specified by *rwlock* does not refer to an initialized read-write lock  
35150 object, or the *abs\_timeout* nanosecond value is less than zero or greater than or  
35151 equal to 1 000 million.

35152 This function shall not return an error code of [EINTR].

## 35153 EXAMPLES

35154 None.

## 35155 APPLICATION USAGE

35156 Applications using this function may be subject to priority inversion, as discussed in the Base  
35157 Definitions volume of IEEE Std 1003.1-200x, Section 3.285, Priority Inversion.

35158 The *pthread\_rwlock\_timedrdlock()* function is part of the Threads and Timeouts options and need  
35159 not be provided on all implementations.

## 35160 RATIONALE

35161 None.

## 35162 FUTURE DIRECTIONS

35163 None.

## 35164 SEE ALSO

35165 *pthread\_rwlock\_destroy()*, *pthread\_rwlock\_init()*, *pthread\_rwlock\_rdlock()*,  
35166 *pthread\_rwlock\_timedwrlock()*, *pthread\_rwlock\_tryrdlock()*, *pthread\_rwlock\_trywrlock()*,  
35167 *pthread\_rwlock\_unlock()*, *pthread\_rwlock\_wrlock()*, the Base Definitions volume of  
35168 IEEE Std 1003.1-200x, <pthread.h>, <time.h>

## 35169 CHANGE HISTORY

35170 First released in Issue 6. Derived from IEEE Std 1003.1j-2000.

35171 **NAME**

35172 pthread\_rwlock\_timedwrlock — lock a read-write lock for writing

35173 **SYNOPSIS**

35174 THR TMO #include &lt;pthread.h&gt;

35175 #include &lt;time.h&gt;

35176 int pthread\_rwlock\_timedwrlock(pthread\_rwlock\_t \*restrict *rwlock*,35177 const struct timespec \*restrict *abs\_timeout*);

35178

35179 **DESCRIPTION**

35180 The *pthread\_rwlock\_timedwrlock()* function shall apply a write lock to the read-write lock  
 35181 referenced by *rwlock* as in the *pthread\_rwlock\_wrlock()* function. However, if the lock cannot be  
 35182 acquired without waiting for other threads to unlock the lock, this wait shall be terminated  
 35183 when the specified timeout expires. The timeout shall expire when the absolute time specified  
 35184 by *abs\_timeout* passes, as measured by the clock on which timeouts are based (that is, when the  
 35185 value of that clock equals or exceeds *abs\_timeout*), or if the absolute time specified by *abs\_timeout*  
 35186 has already been passed at the time of the call.

35187 TMR If the Timers option is supported, the timeout shall be based on the `CLOCK_REALTIME` clock. If  
 35188 the Timers option is not supported, the timeout shall be based on the system clock as returned  
 35189 by the *time()* function. The resolution of the timeout shall be the resolution of the clock on which  
 35190 it is based. The `timespec` data type is defined in the `<time.h>` header. Under no circumstances  
 35191 shall the function fail with a timeout if the lock can be acquired immediately. The validity of the  
 35192 *abs\_timeout* parameter need not be checked if the lock can be immediately acquired.

35193 If a signal that causes a signal handler to be executed is delivered to a thread blocked on a read-  
 35194 write lock via a call to *pthread\_rwlock\_timedwrlock()*, upon return from the signal handler the  
 35195 thread shall resume waiting for the lock as if it was not interrupted.

35196 The calling thread may deadlock if at the time the call is made it holds the read-write lock. The  
 35197 results are undefined if this function is called with an uninitialized read-write lock.

35198 **RETURN VALUE**

35199 The *pthread\_rwlock\_timedwrlock()* function shall return zero if the lock for writing on the read-  
 35200 write lock object referenced by *rwlock* is acquired. Otherwise, an error number shall be returned  
 35201 to indicate the error.

35202 **ERRORS**35203 The *pthread\_rwlock\_timedwrlock()* function shall fail if:

35204 [ETIMEDOUT] The lock could not be acquired before the specified timeout expired.

35205 The *pthread\_rwlock\_timedwrlock()* function may fail if:35206 [EDEADLK] The calling thread already holds the *rwlock*.

35207 [EINVAL] The value specified by *rwlock* does not refer to an initialized read-write lock  
 35208 object, or the *abs\_timeout* nanosecond value is less than zero or greater than or  
 35209 equal to 1 000 million.

35210 This function shall not return an error code of [EINTR].

## 35211 EXAMPLES

35212 None.

## 35213 APPLICATION USAGE

35214 Applications using this function may be subject to priority inversion, as discussed in the Base  
35215 Definitions volume of IEEE Std 1003.1-200x, Section 3.285, Priority Inversion.

35216 The *pthread\_rwlock\_timedwrlock()* function is part of the Threads and Timeouts options and need  
35217 not be provided on all implementations.

## 35218 RATIONALE

35219 None.

## 35220 FUTURE DIRECTIONS

35221 None.

## 35222 SEE ALSO

35223 *pthread\_rwlock\_destroy()*, *pthread\_rwlock\_init()*, *pthread\_rwlock\_rdlock()*,  
35224 *pthread\_rwlock\_timedrdlock()*, *pthread\_rwlock\_tryrdlock()*, *pthread\_rwlock\_trywrlock()*,  
35225 *pthread\_rwlock\_unlock()*, *pthread\_rwlock\_wrlock()*, the Base Definitions volume of  
35226 IEEE Std 1003.1-200x, <pthread.h>, <time.h>

## 35227 CHANGE HISTORY

35228 First released in Issue 6. Derived from IEEE Std 1003.1j-2000.



35229 **NAME**

35230 pthread\_rwlock\_tryrdlock — lock a read-write lock object for reading

35231 **SYNOPSIS**

35232 THR #include &lt;pthread.h&gt;

35233 int pthread\_rwlock\_tryrdlock(pthread\_rwlock\_t \*rwlock);

35234

35235 **DESCRIPTION**35236 Refer to *pthread\_rwlock\_rdlock()*.

35237 **NAME**

35238 pthread\_rwlock\_trywrlock, pthread\_rwlock\_wrlock — lock a read-write lock object for writing

35239 **SYNOPSIS**

35240 THR #include <pthread.h>

35241 int pthread\_rwlock\_trywrlock(pthread\_rwlock\_t \*rwlock);

35242 int pthread\_rwlock\_wrlock(pthread\_rwlock\_t \*rwlock);

35243

35244 **DESCRIPTION**

35245 The *pthread\_rwlock\_trywrlock()* function shall apply a write lock like the *pthread\_rwlock\_wrlock()*  
 35246 function, with the exception that the function shall fail if any thread currently holds *rwlock* (for  
 35247 reading or writing).

35248 The *pthread\_rwlock\_wrlock()* function shall apply a write lock to the read-write lock referenced  
 35249 by *rwlock*. The calling thread acquires the write lock if no other thread (reader or writer) holds  
 35250 the read-write lock *rwlock*. Otherwise, the thread shall block until it can acquire the lock. The  
 35251 calling thread may deadlock if at the time the call is made it holds the read-write lock (whether a  
 35252 read or write lock).

35253 Implementations may favor writers over readers to avoid writer starvation.

35254 Results are undefined if any of these functions are called with an uninitialized read-write lock.

35255 If a signal is delivered to a thread waiting for a read-write lock for writing, upon return from the  
 35256 signal handler the thread resumes waiting for the read-write lock for writing as if it was not  
 35257 interrupted.

35258 **RETURN VALUE**

35259 The *pthread\_rwlock\_trywrlock()* function shall return zero if the lock for writing on the read-write  
 35260 lock object referenced by *rwlock* is acquired. Otherwise, an error number shall be returned to  
 35261 indicate the error.

35262 If successful, the *pthread\_rwlock\_wrlock()* function shall return zero; otherwise, an error number  
 35263 shall be returned to indicate the error.

35264 **ERRORS**

35265 The *pthread\_rwlock\_trywrlock()* function shall fail if:

35266 [EBUSY] The read-write lock could not be acquired for writing because it was already  
 35267 locked for reading or writing.

35268 The *pthread\_rwlock\_trywrlock()* and *pthread\_rwlock\_wrlock()* functions may fail if:

35269 [EINVAL] The value specified by *rwlock* does not refer to an initialized read-write lock  
 35270 object.

35271 The *pthread\_rwlock\_wrlock()* function may fail if:

35272 [EDEADLK] The current thread already owns the read-write lock for writing or reading.

35273 These functions shall not return an error code of [EINTR].

35274 **EXAMPLES**

35275 None.

35276 **APPLICATION USAGE**

35277 Applications using these functions may be subject to priority inversion, as discussed in the Base  
35278 Definitions volume of IEEE Std 1003.1-200x, Section 3.285, Priority Inversion.

35279 **RATIONALE**

35280 None.

35281 **FUTURE DIRECTIONS**

35282 None.

35283 **SEE ALSO**

35284 *pthread\_rwlock\_destroy()*, *pthread\_rwlock\_init()*, *pthread\_rwlock\_rdlock()*,  
35285 *pthread\_rwlock\_timedrdlock()*, *pthread\_rwlock\_timedwrlock()*, *pthread\_rwlock\_tryrdlock()*,  
35286 *pthread\_rwlock\_unlock()*, the Base Definitions volume of IEEE Std 1003.1-200x, <**pthread.h**>

35287 **CHANGE HISTORY**

35288 First released in Issue 5.

35289 **Issue 6**

35290 The following changes are made for alignment with IEEE Std 1003.1j-2000:

- 35291 • The margin code in the SYNOPSIS is changed to THR to indicate that the functionality is  
35292 now part of the Threads option (previously it was part of the Read-Write Locks option in  
35293 IEEE Std 1003.1j-2000 and also part of the XSI extension).
- 35294 • The [EDEADLK] error is deleted as a *pthread\_rwlock\_trywrlock()* error.
- 35295 • The SEE ALSO section is updated.

35296 **NAME**

35297 pthread\_rwlock\_unlock — unlock a read-write lock object

35298 **SYNOPSIS**

35299 THR #include &lt;pthread.h&gt;

35300 int pthread\_rwlock\_unlock(pthread\_rwlock\_t \*rwlock);

35301

35302 **DESCRIPTION**

35303 The *pthread\_rwlock\_unlock()* function shall release a lock held on the read-write lock object |  
35304 referenced by *rwlock*. Results are undefined if the read-write lock *rwlock* is not held by the |  
35305 calling thread.

35306 If this function is called to release a read lock from the read-write lock object and there are other |  
35307 read locks currently held on this read-write lock object, the read-write lock object remains in the |  
35308 read locked state. If this function releases the last read lock for this read-write lock object, the |  
35309 read-write lock object shall be put in the unlocked state with no owners.

35310 If this function is called to release a write lock for this read-write lock object, the read-write lock |  
35311 object shall be put in the unlocked state.

35312 If there are threads blocked on the lock when it becomes available, the scheduling policy shall |  
35313 TPS determine which thread(s) shall acquire the lock. If the Thread Execution Scheduling option is |  
35314 supported, when threads executing with the scheduling policies SCHED\_FIFO, SCHED\_RR, or |  
35315 SCHED\_SPORADIC are waiting on the lock, they shall acquire the lock in priority order when |  
35316 the lock becomes available. For equal priority threads, write locks shall take precedence over |  
35317 read locks. If the Thread Execution Scheduling option is not supported, it is implementation- |  
35318 defined whether write locks take precedence over read locks.

35319 Results are undefined if any of these functions are called with an uninitialized read-write lock.

35320 **RETURN VALUE**

35321 If successful, the *pthread\_rwlock\_unlock()* function shall return zero; otherwise, an error number |  
35322 shall be returned to indicate the error.

35323 **ERRORS**

35324 The *pthread\_rwlock\_unlock()* function may fail if:

35325 [EINVAL] The value specified by *rwlock* does not refer to an initialized read-write lock |  
35326 object.

35327 [EPERM] The current thread does not hold a lock on the read-write lock.

35328 The *pthread\_rwlock\_unlock()* function shall not return an error code of [EINTR].

35329 **EXAMPLES**

35330 None.

35331 **APPLICATION USAGE**

35332 None.

35333 **RATIONALE**

35334 None.

35335 **FUTURE DIRECTIONS**

35336 None.

35337 **SEE ALSO**

35338 *pthread\_rwlock\_destroy()*, *pthread\_rwlock\_init()*, *pthread\_rwlock\_rdlock()*,  
35339 *pthread\_rwlock\_timedrdlock()*, *pthread\_rwlock\_timedwrlock()*, *pthread\_rwlock\_tryrdlock()*,  
35340 *pthread\_rwlock\_trywrlock()*, *pthread\_rwlock\_wrlock()*, the Base Definitions volume of  
35341 IEEE Std 1003.1-200x, <**pthread.h**>

35342 **CHANGE HISTORY**

35343 First released in Issue 5.

35344 **Issue 6**

35345 The following changes are made for alignment with IEEE Std 1003.1j-2000:

- 35346 • The margin code in the SYNOPSIS is changed to THR to indicate that the functionality is  
35347 now part of the Threads option (previously it was part of the Read-Write Locks option in  
35348 IEEE Std 1003.1j-2000 and also part of the XSI extension).
- 35349 • The DESCRIPTION is updated as follows:
  - 35350 — The conditions under which writers have precedence over readers are specified.
  - 35351 — The concept of read-write lock owner is deleted.
- 35352 • The SEE ALSO section is updated.

35353 **NAME**

35354 pthread\_rwlock\_wrlock — lock a read-write lock object for writing

35355 **SYNOPSIS**

35356 THR #include <pthread.h>

35357 int pthread\_rwlock\_wrlock(pthread\_rwlock\_t \**rwlock*);

35358

35359 **DESCRIPTION**

35360 Refer to *pthread\_rwlock\_trywrlock()*.

35361 **NAME**

35362 pthread\_rwlockattr\_destroy, pthread\_rwlockattr\_init — destroy and initialize read-write lock  
 35363 attributes object

35364 **SYNOPSIS**

```
35365 THR #include <pthread.h>
```

```
35366 int pthread_rwlockattr_destroy(pthread_rwlockattr_t *attr);
```

```
35367 int pthread_rwlockattr_init(pthread_rwlockattr_t *attr);
```

```
35368
```

35369 **DESCRIPTION**

35370 The *pthread\_rwlockattr\_destroy()* function shall destroy a read-write lock attributes object. A  
 35371 destroyed *attr* attributes object can be reinitialized using *pthread\_rwlockattr\_init()*; the results of  
 35372 otherwise referencing the object after it has been destroyed are undefined. An implementation  
 35373 may cause *pthread\_rwlockattr\_destroy()* to set the object referenced by *attr* to an invalid value.

35374 The *pthread\_rwlockattr\_init()* function shall initialize a read-write lock attributes object *attr* with  
 35375 the default value for all of the attributes defined by the implementation.

35376 Results are undefined if *pthread\_rwlockattr\_init()* is called specifying an already initialized *attr*  
 35377 attributes object.

35378 After a read-write lock attributes object has been used to initialize one or more read-write locks,  
 35379 any function affecting the attributes object (including destruction) shall not affect any previously  
 35380 initialized read-write locks.

35381 **RETURN VALUE**

35382 If successful, the *pthread\_rwlockattr\_destroy()* and *pthread\_rwlockattr\_init()* functions shall return  
 35383 zero; otherwise, an error number shall be returned to indicate the error.

35384 **ERRORS**

35385 The *pthread\_rwlockattr\_destroy()* function may fail if:

35386 [EINVAL] The value specified by *attr* is invalid.

35387 The *pthread\_rwlockattr\_init()* function shall fail if:

35388 [ENOMEM] Insufficient memory exists to initialize the read-write lock attributes object.

35389 These functions shall not return an error code of [EINTR].

35390 **EXAMPLES**

35391 None.

35392 **APPLICATION USAGE**

35393 None.

35394 **RATIONALE**

35395 None.

35396 **FUTURE DIRECTIONS**

35397 None.

35398 **SEE ALSO**

35399 *pthread\_rwlock\_init()*, *pthread\_rwlockattr\_getpshared()*, *pthread\_rwlockattr\_setpshared()*, the Base  
 35400 Definitions volume of IEEE Std 1003.1-200x, <pthread.h>

35401 **CHANGE HISTORY**

35402 First released in Issue 5.

35403 **Issue 6**

35404 The following changes are made for alignment with IEEE Std 1003.1j-2000:

- 35405 • The margin code in the SYNOPSIS is changed to THR to indicate that the functionality is
- 35406 now part of the Threads option (previously it was part of the Read-Write Locks option in
- 35407 IEEE Std 1003.1j-2000 and also part of the XSI extension).
- 35408 • The SEE ALSO section is updated.



35409 **NAME**

35410 pthread\_rwlockattr\_getpshared, pthread\_rwlockattr\_setpshared — get and set process-shared  
 35411 attribute of read-write lock attributes object

35412 **SYNOPSIS**

```
35413 THR TSH #include <pthread.h>
```

```
35414 int pthread_rwlockattr_getpshared(  
35415     const pthread_rwlockattr_t *restrict attr,  
35416     int *restrict pshared);  
35417 int pthread_rwlockattr_setpshared(pthread_rwlockattr_t *attr,  
35418     int pshared);  
35419
```

35420 **DESCRIPTION**

35421 The *pthread\_rwlockattr\_getpshared()* function shall obtain the value of the *process-shared* attribute  
 35422 from the initialized attributes object referenced by *attr*. The *pthread\_rwlockattr\_setpshared()* |  
 35423 function shall set the *process-shared* attribute in an initialized attributes object referenced by *attr*. |

35424 The *process-shared* attribute shall be set to PTHREAD\_PROCESS\_SHARED to permit a read-  
 35425 write lock to be operated upon by any thread that has access to the memory where the read-  
 35426 write lock is allocated, even if the read-write lock is allocated in memory that is shared by  
 35427 multiple processes. If the *process-shared* attribute is PTHREAD\_PROCESS\_PRIVATE, the read-  
 35428 write lock shall only be operated upon by threads created within the same process as the thread  
 35429 that initialized the read-write lock; if threads of differing processes attempt to operate on such a  
 35430 read-write lock, the behavior is undefined. The default value of the *process-shared* attribute shall  
 35431 be PTHREAD\_PROCESS\_PRIVATE.

35432 Additional attributes, their default values, and the names of the associated functions to get and  
 35433 set those attribute values are implementation-defined.

35434 **RETURN VALUE**

35435 Upon successful completion, the *pthread\_rwlockattr\_getpshared()* shall return zero and store the  
 35436 value of the *process-shared* attribute of *attr* into the object referenced by the *pshared* parameter.  
 35437 Otherwise, an error number shall be returned to indicate the error.

35438 If successful, the *pthread\_rwlockattr\_setpshared()* function shall return zero; otherwise, an error  
 35439 number shall be returned to indicate the error.

35440 **ERRORS**

35441 The *pthread\_rwlockattr\_getpshared()* and *pthread\_rwlockattr\_setpshared()* functions may fail if:

35442 [EINVAL] The value specified by *attr* is invalid.

35443 The *pthread\_rwlockattr\_setpshared()* function may fail if:

35444 [EINVAL] The new value specified for the attribute is outside the range of legal values  
 35445 for that attribute.

35446 These functions shall not return an error code of [EINTR].

35447 **EXAMPLES**

35448 None.

35449 **APPLICATION USAGE**

35450 None.

35451 **RATIONALE**

35452 None.

35453 **FUTURE DIRECTIONS**

35454 None.

35455 **SEE ALSO**

35456 *pthread\_rwlock\_init()*, *pthread\_rwlockattr\_destroy()*, *pthread\_rwlockattr\_init()*, the Base Definitions  
35457 volume of IEEE Std 1003.1-200x, <pthread.h>

35458 **CHANGE HISTORY**

35459 First released in Issue 5.

35460 **Issue 6**

35461 The following changes are made for alignment with IEEE Std 1003.1j-2000:

- 35462 • The margin code in the SYNOPSIS is changed to THR TSH to indicate that the functionality  
35463 is now part of the Threads option (previously it was part of the Read-Write Locks option in  
35464 IEEE Std 1003.1j-2000 and also part of the XSI extension).
- 35465 • The DESCRIPTION notes that additional attributes are implementation-defined.
- 35466 • The SEE ALSO section is updated.

35467 The **restrict** keyword is added to the *pthread\_rwlockattr\_getpshared()* prototype for alignment  
35468 with the ISO/IEC 9899:1999 standard.

35469 **NAME**

35470 pthread\_rwlockattr\_init — initialize read-write lock attributes object

35471 **SYNOPSIS**

35472 XSI #include &lt;pthread.h&gt;

35473 int pthread\_rwlockattr\_init(pthread\_rwlockattr\_t \*attr);

35474

35475 **DESCRIPTION**35476 Refer to *pthread\_rwlockattr\_destroy()*.

35477 **NAME**

35478 pthread\_rwlockattr\_setpshared — set process-shared attribute of read-write lock attributes  
35479 object

35480 **SYNOPSIS**

35481 XSI #include <pthread.h>

```
35482 int pthread_rwlockattr_setpshared(pthread_rwlockattr_t *attr,  
35483 int pshared);
```

35484

35485 **DESCRIPTION**

35486 Refer to *pthread\_rwlockattr\_getpshared()*.

35487 **NAME**

35488 pthread\_self — get calling thread's ID

35489 **SYNOPSIS**

35490 THR #include &lt;pthread.h&gt;

35491 pthread\_t pthread\_self(void);

35492

35493 **DESCRIPTION**35494 The *pthread\_self()* function shall return the thread ID of the calling thread.35495 **RETURN VALUE**

35496 Refer to the DESCRIPTION.

35497 **ERRORS**

35498 No errors are defined.

35499 The *pthread\_self()* function shall not return an error code of [EINTR].35500 **EXAMPLES**

35501 None.

35502 **APPLICATION USAGE**

35503 None.

35504 **RATIONALE**35505 The *pthread\_self()* function provides a capability similar to the *getpid()* function for processes  
35506 and the rationale is the same: the creation call does not provide the thread ID to the created  
35507 thread.35508 **FUTURE DIRECTIONS**

35509 None.

35510 **SEE ALSO**35511 *pthread\_create()*, *pthread\_equal()*, the Base Definitions volume of IEEE Std 1003.1-200x,  
35512 <pthread.h>35513 **CHANGE HISTORY**

35514 First released in Issue 5. Included for alignment with the POSIX Threads Extension.

35515 **Issue 6**35516 The *pthread\_self()* function is marked as part of the Threads option.

35517 **NAME**

35518 pthread\_setcancelstate, pthread\_setcanceltype, pthread\_testcancel — set cancelability state

35519 **SYNOPSIS**

35520 THR #include &lt;pthread.h&gt;

35521 int pthread\_setcancelstate(int state, int \*oldstate);

35522 int pthread\_setcanceltype(int type, int \*oldtype);

35523 void pthread\_testcancel(void);

35524

35525 **DESCRIPTION**

35526 The *pthread\_setcancelstate()* function shall atomically both set the calling thread's cancelability  
35527 state to the indicated *state* and return the previous cancelability state at the location referenced  
35528 by *oldstate*. Legal values for *state* are PTHREAD\_CANCEL\_ENABLE and  
35529 PTHREAD\_CANCEL\_DISABLE.

35530 The *pthread\_setcanceltype()* function shall atomically both set the calling thread's cancelability  
35531 type to the indicated *type* and return the previous cancelability type at the location referenced by  
35532 *oldtype*. Legal values for *type* are PTHREAD\_CANCEL\_DEFERRED and  
35533 PTHREAD\_CANCEL\_ASYNCHRONOUS.

35534 The cancelability state and type of any newly created threads, including the thread in which  
35535 *main()* was first invoked, shall be PTHREAD\_CANCEL\_ENABLE and  
35536 PTHREAD\_CANCEL\_DEFERRED respectively.

35537 The *pthread\_testcancel()* function shall create a cancellation point in the calling thread. The  
35538 *pthread\_testcancel()* function shall have no effect if cancelability is disabled.

35539 **RETURN VALUE**

35540 If successful, the *pthread\_setcancelstate()* and *pthread\_setcanceltype()* functions shall return zero;  
35541 otherwise, an error number shall be returned to indicate the error.

35542 **ERRORS**

35543 The *pthread\_setcancelstate()* function may fail if:

35544 [EINVAL] The specified state is not PTHREAD\_CANCEL\_ENABLE or  
35545 PTHREAD\_CANCEL\_DISABLE.

35546 The *pthread\_setcanceltype()* function may fail if:

35547 [EINVAL] The specified type is not PTHREAD\_CANCEL\_DEFERRED or  
35548 PTHREAD\_CANCEL\_ASYNCHRONOUS.

35549 These functions shall not return an error code of [EINTR].

35550 **EXAMPLES**

35551 None.

35552 **APPLICATION USAGE**

35553 None.

35554 **RATIONALE**

35555 The *pthread\_setcancelstate()* and *pthread\_setcanceltype()* functions control the points at which a |  
35556 thread may be asynchronously canceled. For cancellation control to be usable in modular fashion, |  
35557 some rules need to be followed.

35558 An object can be considered to be a generalization of a procedure. It is a set of procedures and  
35559 global variables written as a unit and called by clients not known by the object. Objects may  
35560 depend on other objects.

35561 First, cancelability should only be disabled on entry to an object, never explicitly enabled. On  
35562 exit from an object, the cancelability state should always be restored to its value on entry to the  
35563 object.

35564 This follows from a modularity argument: if the client of an object (or the client of an object that  
35565 uses that object) has disabled cancelability, it is because the client does not want to be concerned  
35566 about cleaning up if the thread is canceled while executing some sequence of actions. If an object  
35567 is called in such a state and it enables cancelability and a cancelation request is pending for that  
35568 thread, then the thread is canceled, contrary to the wish of the client that disabled.

35569 Second, the cancelability type may be explicitly set to either *deferred* or *asynchronous* upon entry  
35570 to an object. But as with the cancelability state, on exit from an object the cancelability type  
35571 should always be restored to its value on entry to the object.

35572 Finally, only functions that are cancel-safe may be called from a thread that is asynchronously  
35573 cancelable.

#### 35574 **FUTURE DIRECTIONS**

35575 None.

#### 35576 **SEE ALSO**

35577 *pthread\_cancel()*, the Base Definitions volume of IEEE Std 1003.1-200x, <pthread.h>

#### 35578 **CHANGE HISTORY**

35579 First released in Issue 5. Included for alignment with the POSIX Threads Extension.

#### 35580 **Issue 6**

35581 The *pthread\_setcancelstate()*, *pthread\_setcanceltype()*, and *pthread\_testcancel()* functions are marked  
35582 as part of the Threads option.

35583 **NAME**

35584 pthread\_setconcurrency — set level of concurrency

35585 **SYNOPSIS**

35586 XSI #include <pthread.h>

35587 int pthread\_setconcurrency(int new\_level);

35588

35589 **DESCRIPTION**

35590 Refer to *pthread\_getconcurrency()*.



35591 **NAME**

35592 pthread\_setschedparam — dynamic thread scheduling parameters access (**REALTIME**  
35593 **THREADS**)

35594 **SYNOPSIS**

```
35595 THR TPS #include <pthread.h>
```

```
35596 int pthread_setschedparam(pthread_t thread, int policy,  
35597     const struct sched_param *param);
```

35598

35599 **DESCRIPTION**

35600 Refer to *pthread\_getschedparam()*.

## 35601 NAME

35602 pthread\_setschedprio — dynamic thread scheduling parameters access (**REALTIME**  
35603 **THREADS**)

## 35604 SYNOPSIS

```
35605 THR TPS #include <pthread.h>
```

```
35606 int pthread_setschedprio(pthread_t thread, int prio);
```

35607

## 35608 DESCRIPTION

35609 The *pthread\_setschedprio()* function shall set the scheduling priority for the thread whose thread  
35610 ID is given by *thread* to the value given by *prio*. See **Scheduling Policies** (on page 494) for a  
35611 description on how this function call affects the ordering of the thread in the thread list for its  
35612 new priority.

35613 If the *pthread\_setschedprio()* function fails, the scheduling priority of the target thread shall not be  
35614 changed.

## 35615 RETURN VALUE

35616 If successful, the *pthread\_setschedprio()* function shall return zero; otherwise, an error number  
35617 shall be returned to indicate the error.

## 35618 ERRORS

35619 The *pthread\_setschedprio()* function may fail if:

35620 [EINVAL] The value of *prio* is invalid for the scheduling policy of the specified thread.

35621 [ENOTSUP] An attempt was made to set the priority to an unsupported value.

35622 [EPERM] The caller does not have the appropriate permission to set the scheduling  
35623 policy of the specified thread.

35624 [EPERM] The implementation does not allow the application to modify the priority to  
35625 the value specified.

35626 [ESRCH] The value specified by *thread* does not refer to an existing thread.

35627 The *pthread\_setschedprio()* function shall not return an error code of [EINTR].

## 35628 EXAMPLES

35629 None.

## 35630 APPLICATION USAGE

35631 None.

## 35632 RATIONALE

35633 The *pthread\_setschedprio()* function provides a way for an application to temporarily raise its  
35634 priority and then lower it again, without having the undesired side effect of yielding to other  
35635 threads of the same priority. This is necessary if the application is to implement its own  
35636 strategies for bounding priority inversion, such as priority inheritance or priority ceilings. This  
35637 capability is especially important if the implementation does not support the Thread Priority  
35638 Protection or Thread Priority Inheritance options, but even if those options are supported it is  
35639 needed if the application is to bound priority inheritance for other resources, such as  
35640 semaphores.

35641 The standard developers considered that while it might be preferable conceptually to solve this  
35642 problem by modifying the specification of *pthread\_setschedparam()*, it was too late to make such a  
35643 change, as there may be implementations that would need to be changed. Therefore, this new  
35644 function was introduced.

35645 **FUTURE DIRECTIONS**

35646       None.

35647 **SEE ALSO**

35648       *pthread\_getschedparam()*, the Base Definitions volume of IEEE Std 1003.1-200x, <pthread.h>

35649 **CHANGE HISTORY**

35650       First released in Issue 6. Included as a response to IEEE PASC Interpretation 1003.1 #96.

35651 **NAME**

35652 pthread\_setspecific — thread-specific data management

35653 **SYNOPSIS**

35654 THR #include <pthread.h>

35655 int pthread\_setspecific(pthread\_key\_t key, const void \*value);

35656

35657 **DESCRIPTION**

35658 Refer to *pthread\_getspecific()*.

35659 **NAME**

35660 pthread\_sigmask, sigprocmask — examine and change blocked signals

35661 **SYNOPSIS**

35662 #include &lt;signal.h&gt;

35663 THR int pthread\_sigmask(int how, const sigset\_t \*restrict set,  
35664 sigset\_t \*restrict oset);35665 CX int sigprocmask(int how, const sigset\_t \*restrict set,  
35666 sigset\_t \*restrict oset);

35667

35668 **DESCRIPTION**35669 THR The *pthread\_sigmask()* function shall examine or change (or both) the calling thread's signal  
35670 mask, regardless of the number of threads in the process. The function shall be equivalent to  
35671 *sigprocmask()*, without the restriction that the call be made in a single-threaded process.35672 In a single-threaded process, the *sigprocmask()* function shall examine or change (or both) the  
35673 signal mask of the calling thread.35674 If the argument *set* is not a null pointer, it points to a set of signals to be used to change the  
35675 currently blocked set.35676 The argument *how* indicates the way in which the set is changed, and the application shall  
35677 ensure it consists of one of the following values:35678 SIG\_BLOCK The resulting set shall be the union of the current set and the signal set  
35679 pointed to by *set*.35680 SIG\_SETMASK The resulting set shall be the signal set pointed to by *set*.35681 SIG\_UNBLOCK The resulting set shall be the intersection of the current set and the  
35682 complement of the signal set pointed to by *set*.35683 If the argument *oset* is not a null pointer, the previous mask shall be stored in the location  
35684 pointed to by *oset*. If *set* is a null pointer, the value of the argument *how* is not significant and the  
35685 process' signal mask shall be unchanged; thus the call can be used to enquire about currently  
35686 blocked signals.35687 If there are any pending unblocked signals after the call to *sigprocmask()*, at least one of those  
35688 signals shall be delivered before the call to *sigprocmask()* returns.35689 It is not possible to block those signals which cannot be ignored. This shall be enforced by the  
35690 system without causing an error to be indicated.35691 If any of the SIGFPE, SIGILL, SIGSEGV, or SIGBUS signals are generated while they are blocked,  
35692 the result is undefined, unless the signal was generated by the *kill()* function, the *sigqueue()*  
35693 function, or the *raise()* function.35694 If *sigprocmask()* fails, the thread's signal mask shall not be changed.35695 The use of the *sigprocmask()* function is unspecified in a multi-threaded process.35696 **RETURN VALUE**35697 THR Upon successful completion *pthread\_sigmask()* shall return 0; otherwise, it shall return the  
35698 corresponding error number.35699 Upon successful completion, *sigprocmask()* shall return 0; otherwise, -1 shall be returned, *errno*  
35700 shall be set to indicate the error, and the process' signal mask shall be unchanged.

35701 **ERRORS**

35702 THR The `pthread_sigmask()` and `sigprocmask()` functions shall fail if:

35703 [EINVAL] The value of the *how* argument is not equal to one of the defined values.

35704 THR The `pthread_sigmask()` function shall not return an error code of [EINTR].

35705 **EXAMPLES**

35706 None.

35707 **APPLICATION USAGE**

35708 None.

35709 **RATIONALE**

35710 When a process' signal mask is changed in a signal-catching function that is installed by  
 35711 `sigaction()`, the restoration of the signal mask on return from the signal-catching function  
 35712 overrides that change (see `sigaction()`). If the signal-catching function was installed with  
 35713 `signal()`, it is unspecified whether this occurs.

35714 See `kill()` for a discussion of the requirement on delivery of signals.

35715 **FUTURE DIRECTIONS**

35716 None.

35717 **SEE ALSO**

35718 `sigaction()`, `sigaddset()`, `sigdelset()`, `sigemptyset()`, `sigfillset()`, `sigismember()`, `sigpending()`,  
 35719 `sigqueue()`, `sigsuspend()`, the Base Definitions volume of IEEE Std 1003.1-200x, <**signal.h**>

35720 **CHANGE HISTORY**

35721 First released in Issue 3.

35722 Entry included for alignment with the POSIX.1-1988 standard.

35723 **Issue 5**

35724 The DESCRIPTION is updated for alignment with the POSIX Threads Extension.

35725 The `pthread_sigmask()` function is added for alignment with the POSIX Threads Extension.

35726 **Issue 6**

35727 The `pthread_sigmask()` function is marked as part of the Threads option.

35728 The SYNOPSIS for `sigprocmask()` is marked as a CX extension to note that the presence of this  
 35729 function in the <**signal.h**> header is an extension to the ISO C standard. |

35730 The following changes are made for alignment with the ISO POSIX-1: 1996 standard: |

35731 • The DESCRIPTION is updated to explicitly state the functions which may generate the  
 35732 signal.

35733 The DESCRIPTION is updated to avoid use of the term “must” for application requirements.

35734 The **restrict** keyword is added to the `pthread_sigmask()` and `sigprocmask()` prototypes for  
 35735 alignment with the ISO/IEC 9899: 1999 standard.

35736 **NAME**

35737 pthread\_spin\_destroy, pthread\_spin\_init — destroy or initialize a spin lock object (**ADVANCED**  
 35738 **REALTIME THREADS**)

35739 **SYNOPSIS**

35740 THR SPI #include <pthread.h>

35741 int pthread\_spin\_destroy(pthread\_spinlock\_t \*lock);

35742 int pthread\_spin\_init(pthread\_spinlock\_t \*lock, int pshared);

35743

35744 **DESCRIPTION**

35745 The *pthread\_spin\_destroy()* function shall destroy the spin lock referenced by *lock* and release any  
 35746 resources used by the lock. The effect of subsequent use of the lock is undefined until the lock is  
 35747 reinitialized by another call to *pthread\_spin\_init()*. The results are undefined if  
 35748 *pthread\_spin\_destroy()* is called when a thread holds the lock, or if this function is called with an  
 35749 uninitialized thread spin lock.

35750 The *pthread\_spin\_init()* function shall allocate any resources required to use the spin lock  
 35751 referenced by *lock* and initialize the lock to an unlocked state.

35752 TSH If the Thread Process-Shared Synchronization option is supported and the value of *pshared* is  
 35753 PTHREAD\_PROCESS\_SHARED, the implementation shall permit the spin lock to be operated  
 35754 upon by any thread that has access to the memory where the spin lock is allocated, even if it is  
 35755 allocated in memory that is shared by multiple processes.

35756 If the Thread Process-Shared Synchronization option is supported and the value of *pshared* is  
 35757 PTHREAD\_PROCESS\_PRIVATE, or if the option is not supported, the spin lock shall only be  
 35758 operated upon by threads created within the same process as the thread that initialized the spin  
 35759 lock. If threads of differing processes attempt to operate on such a spin lock, the behavior is  
 35760 undefined.

35761 The results are undefined if *pthread\_spin\_init()* is called specifying an already initialized spin  
 35762 lock. The results are undefined if a spin lock is used without first being initialized.

35763 If the *pthread\_spin\_init()* function fails, the lock is not initialized and the contents of *lock* are  
 35764 undefined.

35765 Only the object referenced by *lock* may be used for performing synchronization.

35766 The result of referring to copies of that object in calls to *pthread\_spin\_destroy()*,  
 35767 *pthread\_spin\_lock()*, *pthread\_spin\_trylock()*, or *pthread\_spin\_unlock()* is undefined.

35768 **RETURN VALUE**

35769 Upon successful completion, these functions shall return zero; otherwise, an error number shall  
 35770 be returned to indicate the error.

35771 **ERRORS**

35772 These functions may fail if:

35773 [EBUSY] The implementation has detected an attempt to initialize or destroy a spin  
 35774 lock while it is in use (for example, while being used in a *pthread\_spin\_lock()*  
 35775 call) by another thread.

35776 [EINVAL] The value specified by *lock* is invalid.

35777 The *pthread\_spin\_init()* function shall fail if:

35778 [EAGAIN] The system lacks the necessary resources to initialize another spin lock.

35779 [ENOMEM] Insufficient memory exists to initialize the lock.

35780 These functions shall not return an error code of [EINTR].

## 35781 EXAMPLES

35782 None.

## 35783 APPLICATION USAGE

35784 The *pthread\_spin\_destroy()* and *pthread\_spin\_init()* functions are part of the Spin Locks option  
35785 and need not be provided on all implementations.

## 35786 RATIONALE

35787 None.

## 35788 FUTURE DIRECTIONS

35789 None.

## 35790 SEE ALSO

35791 *pthread\_spin\_lock()*, *pthread\_spin\_trylock()*, *pthread\_spin\_unlock()*, the Base Definitions volume of  
35792 IEEE Std 1003.1-200x, <<pthread.h>>

## 35793 CHANGE HISTORY

35794 First released in Issue 6. Derived from IEEE Std 1003.1j-2000.

35795 In the SYNOPSIS, the inclusion of <sys/types.h> is no longer required.



35796 **NAME**

35797 pthread\_spin\_init — initialize a spin lock object (**ADVANCED REALTIME THREADS**)

35798 **SYNOPSIS**

35799 THR SPI #include <pthread.h>

35800 int pthread\_spin\_init(pthread\_spinlock\_t \*lock, int pshared);

35801

35802 **DESCRIPTION**

35803 Refer to *pthread\_spin\_destroy()*.

## 35804 NAME

35805 pthread\_spin\_lock, pthread\_spin\_trylock — lock a spin lock object (ADVANCED REALTIME  
35806 THREADS)

## 35807 SYNOPSIS

```
35808 THR SPI #include <pthread.h>
```

```
35809 int pthread_spin_lock(pthread_spinlock_t *lock);  
35810 int pthread_spin_trylock(pthread_spinlock_t *lock);  
35811
```

## 35812 DESCRIPTION

35813 The *pthread\_spin\_lock()* function shall lock the spin lock referenced by *lock*. The calling thread  
35814 shall acquire the lock if it is not held by another thread. Otherwise, the thread shall spin (that is, |  
35815 shall not return from the *pthread\_spin\_lock()* call) until the lock becomes available. The results |  
35816 are undefined if the calling thread holds the lock at the time the call is made. The  
35817 *pthread\_spin\_trylock()* function shall lock the spin lock referenced by *lock* if it is not held by any |  
35818 thread. Otherwise, the function shall fail. |

35819 The results are undefined if any of these functions is called with an uninitialized spin lock.

## 35820 RETURN VALUE

35821 Upon successful completion, these functions shall return zero; otherwise, an error number shall  
35822 be returned to indicate the error.

## 35823 ERRORS

35824 These functions may fail if:

35825 [EINVAL] The value specified by *lock* does not refer to an initialized spin lock object.

35826 The *pthread\_spin\_lock()* function may fail if:

35827 [EDEADLK] The calling thread already holds the lock.

35828 The *pthread\_spin\_trylock()* function shall fail if:

35829 [EBUSY] A thread currently holds the lock.

35830 These functions shall not return an error code of [EINTR].

## 35831 EXAMPLES

35832 None.

## 35833 APPLICATION USAGE

35834 Applications using this function may be subject to priority inversion, as discussed in the Base  
35835 Definitions volume of IEEE Std 1003.1-200x, Section 3.285, Priority Inversion.

35836 The *pthread\_spin\_lock()* and *pthread\_spin\_trylock()* functions are part of the Spin Locks option  
35837 and need not be provided on all implementations.

## 35838 RATIONALE

35839 None.

## 35840 FUTURE DIRECTIONS

35841 None.

## 35842 SEE ALSO

35843 *pthread\_spin\_init()*, *pthread\_spin\_destroy()*, *pthread\_spin\_unlock()*, the Base Definitions volume of  
35844 IEEE Std 1003.1-200x, <pthread.h>

35845 **CHANGE HISTORY**

35846 First released in Issue 6. Derived from IEEE Std 1003.1j-2000.

35847 In the SYNOPSIS, the inclusion of `<sys/types.h>` is no longer required.

35848 **NAME**

35849 pthread\_spin\_trylock — lock a spin lock object (**ADVANCED REALTIME THREADS**)

35850 **SYNOPSIS**

35851 THR SPI #include <pthread.h>

35852 int pthread\_spin\_trylock(pthread\_spinlock\_t \*lock);

35853

35854 **DESCRIPTION**

35855 Refer to *pthread\_spin\_lock()*.

35856 **NAME**35857 pthread\_spin\_unlock — unlock a spin lock object (**ADVANCED REALTIME THREADS**)35858 **SYNOPSIS**

35859 THR SPI #include &lt;pthread.h&gt;

35860 int pthread\_spin\_unlock(pthread\_spinlock\_t \*lock);

35861

35862 **DESCRIPTION**

35863 The *pthread\_spin\_unlock()* function shall release the spin lock referenced by *lock* which was  
35864 locked via the *pthread\_spin\_lock()* or *pthread\_spin\_trylock()* functions. The results are undefined if  
35865 the lock is not held by the calling thread. If there are threads spinning on the lock when  
35866 *pthread\_spin\_unlock()* is called, the lock becomes available and an unspecified spinning thread  
35867 shall acquire the lock.

35868 The results are undefined if this function is called with an uninitialized thread spin lock.

35869 **RETURN VALUE**

35870 Upon successful completion, the *pthread\_spin\_unlock()* function shall return zero; otherwise, an  
35871 error number shall be returned to indicate the error.

35872 **ERRORS**35873 The *pthread\_spin\_unlock()* function may fail if:

35874 [EINVAL] An invalid argument was specified.

35875 [EPERM] The calling thread does not hold the lock.

35876 This function shall not return an error code of [EINTR].

35877 **EXAMPLES**

35878 None.

35879 **APPLICATION USAGE**

35880 The *pthread\_spin\_unlock()* function is part of the Spin Locks option and need not be provided on  
35881 all implementations.

35882 **RATIONALE**

35883 None.

35884 **FUTURE DIRECTIONS**

35885 None.

35886 **SEE ALSO**

35887 *pthread\_spin\_init()*, *pthread\_spin\_destroy()*, *pthread\_spin\_lock()*, *pthread\_spin\_trylock()*, the Base  
35888 Definitions volume of IEEE Std 1003.1-200x, <pthread.h>

35889 **CHANGE HISTORY**

35890 First released in Issue 6. Derived from IEEE Std 1003.1j-2000.

35891 In the SYNOPSIS, the inclusion of &lt;sys/types.h&gt; is no longer required.

35892 **NAME**

35893 pthread\_testcancel — set cancelability state

35894 **SYNOPSIS**

35895 THR #include <pthread.h>

35896 void pthread\_testcancel(void);

35897

35898 **DESCRIPTION**

35899 Refer to *pthread\_setcancelstate()*.

35900 **NAME**

35901 ptsname — get name of the slave pseudo-terminal device

35902 **SYNOPSIS**

35903 XSI `#include <stdlib.h>`

35904 `char *ptsname(int fildev);`

35905

35906 **DESCRIPTION**

35907 The *ptsname()* function shall return the name of the slave pseudo-terminal device associated  
35908 with a master pseudo-terminal device. The *fildev* argument is a file descriptor that refers to the  
35909 master device. The *ptsname()* function shall return a pointer to a string containing the pathname  
35910 of the corresponding slave device.

35911 The *ptsname()* function need not be reentrant. A function that is not required to be reentrant is  
35912 not required to be thread-safe.

35913 **RETURN VALUE**

35914 Upon successful completion, *ptsname()* shall return a pointer to a string which is the name of the  
35915 pseudo-terminal slave device. Upon failure, *ptsname()* shall return a null pointer. This could  
35916 occur if *fildev* is an invalid file descriptor or if the slave device name does not exist in the file  
35917 system.

35918 **ERRORS**

35919 No errors are defined.

35920 **EXAMPLES**

35921 None.

35922 **APPLICATION USAGE**

35923 The value returned may point to a static data area that is overwritten by each call to *ptsname()*.

35924 **RATIONALE**

35925 None.

35926 **FUTURE DIRECTIONS**

35927 None.

35928 **SEE ALSO**

35929 *grantpt()*, *open()*, *ttyname()*, *unlockpt()*, the Base Definitions volume of IEEE Std 1003.1-200x,  
35930 `<stdlib.h>`

35931 **CHANGE HISTORY**

35932 First released in Issue 4, Version 2.

35933 **Issue 5**

35934 Moved from X/OPEN UNIX extension to BASE.

35935 A note indicating that this function need not be reentrant is added to the DESCRIPTION.

35936 **NAME**

35937       putc — put byte on a stream

35938 **SYNOPSIS**

35939       #include &lt;stdio.h&gt;

35940       int putc(int *c*, FILE \**stream*);35941 **DESCRIPTION**

35942 **CX**       The functionality described on this reference page is aligned with the ISO C standard. Any  
35943 conflict between the requirements described here and the ISO C standard is unintentional. This  
35944 volume of IEEE Std 1003.1-200x defers to the ISO C standard.

35945       The *putc()* function shall be equivalent to *fputc()*, except that if it is implemented as a macro it  
35946 may evaluate *stream* more than once, so the argument should never be an expression with side  
35947 effects.

35948 **RETURN VALUE**35949       Refer to *fputc()*.35950 **ERRORS**35951       Refer to *fputc()*.35952 **EXAMPLES**

35953       None.

35954 **APPLICATION USAGE**

35955       Since it may be implemented as a macro, *putc()* may treat a *stream* argument with side effects |  
35956 incorrectly. In particular, *putc(c,\*f++)* does not necessarily work correctly. Therefore, use of this  
35957 function is not recommended in such situations; *fputc()* should be used instead.

35958 **RATIONALE**

35959       None.

35960 **FUTURE DIRECTIONS**

35961       None.

35962 **SEE ALSO**35963       *fputc()*, the Base Definitions volume of IEEE Std 1003.1-200x, <stdio.h>35964 **CHANGE HISTORY**

35965       First released in Issue 1. Derived from Issue 1 of the SVID.



35966 **NAME**

35967       putc\_unlocked — stdio with explicit client locking

35968 **SYNOPSIS**

35969 TSF     #include &lt;stdio.h&gt;

35970       int putc\_unlocked(int *c*, FILE \**stream*);

35971

35972 **DESCRIPTION**35973       Refer to *getc\_unlocked()*.

35974 **NAME**

35975 putchar — put byte on stdout stream

35976 **SYNOPSIS**

35977 #include <stdio.h>

35978 int putchar(int c);

35979 **DESCRIPTION**

35980 cx The functionality described on this reference page is aligned with the ISO C standard. Any  
35981 conflict between the requirements described here and the ISO C standard is unintentional. This  
35982 volume of IEEE Std 1003.1-200x defers to the ISO C standard.

35983 The function call *putchar(c)* shall be equivalent to *putc(c,stdout)*.

35984 **RETURN VALUE**

35985 Refer to *fputc()*.

35986 **ERRORS**

35987 Refer to *fputc()*.

35988 **EXAMPLES**

35989 None.

35990 **APPLICATION USAGE**

35991 None.

35992 **RATIONALE**

35993 None.

35994 **FUTURE DIRECTIONS**

35995 None.

35996 **SEE ALSO**

35997 *putc()*, the Base Definitions volume of IEEE Std 1003.1-200x, <stdio.h>

35998 **CHANGE HISTORY**

35999 First released in Issue 1. Derived from Issue 1 of the SVID.

36000 **NAME**

36001 putchar\_unlocked — stdio with explicit client locking

36002 **SYNOPSIS**

36003 TSF #include &lt;stdio.h&gt;

36004 int putchar\_unlocked(int c);

36005

36006 **DESCRIPTION**36007 Refer to *getc\_unlocked()*.

36008 **NAME**

36009 putenv — change or add a value to environment

36010 **SYNOPSIS**36011 XSI `#include <stdlib.h>`36012 `int putenv(char *string);`

36013

36014 **DESCRIPTION**

36015 The *putenv()* function shall use the *string* argument to set environment variable values. The  
 36016 *string* argument should point to a string of the form "*name=value*". The *putenv()* function shall  
 36017 make the value of the environment variable *name* equal to *value* by altering an existing variable  
 36018 or creating a new one. In either case, the string pointed to by *string* shall become part of the  
 36019 environment, so altering the string shall change the environment. The space used by *string* is no  
 36020 longer used once a new string-defining *name* is passed to *putenv()*.

36021 The *putenv()* function need not be reentrant. A function that is not required to be reentrant is not  
 36022 required to be thread-safe.

36023 **RETURN VALUE**

36024 Upon successful completion, *putenv()* shall return 0; otherwise, it shall return a non-zero value  
 36025 and set *errno* to indicate the error.

36026 **ERRORS**36027 The *putenv()* function may fail if:

36028 [ENOMEM] Insufficient memory was available.

36029 **EXAMPLES**36030 **Changing the Value of an Environment Variable**

36031 The following example changes the value of the *HOME* environment variable to the value  
 36032 */usr/home*.

36033 `#include <stdlib.h>`36034 `...`36035 `static char *var = "HOME=/usr/home";`36036 `int ret;`36037 `ret = putenv(var);`36038 **APPLICATION USAGE**

36039 The *putenv()* function manipulates the environment pointed to by *environ*, and can be used in  
 36040 conjunction with *getenv()*.

36041 This routine may use *malloc()* to enlarge the environment.

36042 A potential error is to call *putenv()* with an automatic variable as the argument, then return from  
 36043 the calling function while *string* is still part of the environment.

36044 The *setenv()* function is preferred over this function.36045 **RATIONALE**

36046 The standard developers noted that *putenv()* is the only function available to add to the  
 36047 environment without permitting memory leaks.

36048 **FUTURE DIRECTIONS**

36049 None.

36050 **SEE ALSO**36051 *exec*, *getenv()*, *malloc()*, *setenv()*, the Base Definitions volume of IEEE Std 1003.1-200x, <stdlib.h>36052 **CHANGE HISTORY**

36053 First released in Issue 1. Derived from Issue 1 of the SVID.

36054 **Issue 5**36055 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.

36057 A note indicating that this function need not be reentrant is added to the DESCRIPTION.

## 36058 NAME

36059 putmsg, putpmsg — send a message on a STREAM (STREAMS)

## 36060 SYNOPSIS

```
36061 XSR #include <stropts.h>
36062
36062 int putmsg(int fildes, const struct strbuf *ctlptr,
36063           const struct strbuf *dataptr, int flags);
36064 int putpmsg(int fildes, const struct strbuf *ctlptr,
36065            const struct strbuf *dataptr, int band, int flags);
36066
```

## 36067 DESCRIPTION

36068 The *putmsg()* function shall create a message from a process buffer(s) and send the message to a  
 36069 STREAMS file. The message may contain either a data part, a control part, or both. The data and  
 36070 control parts are distinguished by placement in separate buffers, as described below. The  
 36071 semantics of each part are defined by the STREAMS module that receives the message.

36072 The *putpmsg()* function is equivalent to *putmsg()*, except that the process can send messages in  
 36073 different priority bands. Except where noted, all requirements on *putmsg()* also pertain to  
 36074 *putpmsg()*.

36075 The *fildes* argument specifies a file descriptor referencing an open STREAM. The *ctlptr* and  
 36076 *dataptr* arguments each point to a **strbuf** structure.

36077 The *ctlptr* argument points to the structure describing the control part, if any, to be included in  
 36078 the message. The *buf* member in the **strbuf** structure points to the buffer where the control  
 36079 information resides, and the *len* member indicates the number of bytes to be sent. The *maxlen*  
 36080 member is not used by *putmsg()*. In a similar manner, the argument *dataptr* specifies the data, if  
 36081 any, to be included in the message. The *flags* argument indicates what type of message should be  
 36082 sent and is described further below.

36083 To send the data part of a message, the application shall ensure that *dataptr* is not a null pointer  
 36084 and the *len* member of *dataptr* is 0 or greater. To send the control part of a message, the  
 36085 application shall ensure that the corresponding values are set for *ctlptr*. No data (control) part  
 36086 shall be sent if either *dataptr(ctlptr)* is a null pointer or the *len* member of *dataptr(ctlptr)* is set to  
 36087 -1.

36088 For *putmsg()*, if a control part is specified and *flags* is set to RS\_HIPRI, a high priority message  
 36089 shall be sent. If no control part is specified, and *flags* is set to RS\_HIPRI, *putmsg()* shall fail and  
 36090 set *errno* to [EINVAL]. If *flags* is set to 0, a normal message (priority band equal to 0) shall be  
 36091 sent. If a control part and data part are not specified and *flags* is set to 0, no message shall be  
 36092 sent and 0 shall be returned.

36093 For *putpmsg()*, the flags are different. The *flags* argument is a bitmask with the following  
 36094 mutually-exclusive flags defined: MSG\_HIPRI and MSG\_BAND. If *flags* is set to 0, *putpmsg()*  
 36095 shall fail and set *errno* to [EINVAL]. If a control part is specified and *flags* is set to MSG\_HIPRI  
 36096 and *band* is set to 0, a high-priority message shall be sent. If *flags* is set to MSG\_HIPRI and either  
 36097 no control part is specified or *band* is set to a non-zero value, *putpmsg()* shall fail and set *errno* to  
 36098 [EINVAL]. If *flags* is set to MSG\_BAND, then a message shall be sent in the priority band  
 36099 specified by *band*. If a control part and data part are not specified and *flags* is set to MSG\_BAND,  
 36100 no message shall be sent and 0 shall be returned.

36101 The *putmsg()* function shall block if the STREAM write queue is full due to internal flow control  
 36102 conditions, with the following exceptions:

- 36103 • For high-priority messages, *putmsg()* shall not block on this condition and continues  
 36104 processing the message.

36105           • For other messages, *putmsg()* shall not block but shall fail when the write queue is full and  
36106            O\_NONBLOCK is set. |

36107           The *putmsg()* function shall also block, unless prevented by lack of internal resources, while |  
36108            waiting for the availability of message blocks in the STREAM, regardless of priority or whether |  
36109            O\_NONBLOCK has been specified. No partial message shall be sent. |

#### 36110 RETURN VALUE

36111           Upon successful completion, *putmsg()* and *putpmsg()* shall return 0; otherwise, they shall return  
36112            -1 and set *errno* to indicate the error.

#### 36113 ERRORS

36114           The *putmsg()* and *putpmsg()* functions shall fail if:

36115           [EAGAIN]        A non-priority message was specified, the O\_NONBLOCK flag is set, and the  
36116                            STREAM write queue is full due to internal flow control conditions; or buffers  
36117                            could not be allocated for the message that was to be created.

36118           [EBADF]        *fildev* is not a valid file descriptor open for writing.

36119           [EINTR]        A signal was caught during *putmsg()*.

36120           [EINVAL]        An undefined value is specified in *flags*, or *flags* is set to RS\_HIPRI or  
36121                            MSG\_HIPRI and no control part is supplied, or the STREAM or multiplexer  
36122                            referenced by *fildev* is linked (directly or indirectly) downstream from a  
36123                            multiplexer, or *flags* is set to MSG\_HIPRI and *band* is non-zero (for *putpmsg()*  
36124                            only).

36125           [ENOSR]        Buffers could not be allocated for the message that was to be created due to  
36126                            insufficient STREAMS memory resources.

36127           [ENOSTR]        A STREAM is not associated with *fildev*.

36128           [ENXIO]        A hangup condition was generated downstream for the specified STREAM.

36129           [EPIPE] or [EIO] The *fildev* argument refers to a STREAMS-based pipe and the other end of the  
36130                            pipe is closed. A SIGPIPE signal is generated for the calling thread.

36131           [ERANGE]        The size of the data part of the message does not fall within the range  
36132                            specified by the maximum and minimum packet sizes of the topmost  
36133                            STREAM module. This value is also returned if the control part of the message  
36134                            is larger than the maximum configured size of the control part of a message,  
36135                            or if the data part of a message is larger than the maximum configured size of  
36136                            the data part of a message.

36137           In addition, *putmsg()* and *putpmsg()* shall fail if the STREAM head had processed an  
36138            asynchronous error before the call. In this case, the value of *errno* does not reflect the result of  
36139            *putmsg()* or *putpmsg()*, but reflects the prior error. |

## 36140 EXAMPLES

36141 **Sending a High-Priority Message**

36142 The value of *fd* is assumed to refer to an open STREAMS file. This call to *putmsg()* does the  
36143 following:

- 36144 1. Creates a high-priority message with a control part and a data part, using the buffers  
36145 pointed to by *ctrlbuf* and *databuf*, respectively.
- 36146 2. Sends the message to the STREAMS file identified by *fd*.

```
36147 #include <stropts.h>
36148 #include <string.h>
36149 ...
36150 int fd;
36151 char *ctrlbuf = "This is the control part";
36152 char *databuf = "This is the data part";
36153 struct strbuf ctrl;
36154 struct strbuf data;
36155 int ret;

36156 ctrl.buf = ctrlbuf;
36157 ctrl.len = strlen(ctrlbuf);

36158 data.buf = databuf;
36159 data.len = strlen(databuf);

36160 ret = putmsg(fd, &ctrl, &data, MSG_HIPRI);
```

36161 **Using putpmsg()**

36162 This example has the same effect as the previous example. In this example, however, the  
36163 *putpmsg()* function creates and sends the message to the STREAMS file.

```
36164 #include <stropts.h>
36165 #include <string.h>
36166 ...
36167 int fd;
36168 char *ctrlbuf = "This is the control part";
36169 char *databuf = "This is the data part";
36170 struct strbuf ctrl;
36171 struct strbuf data;
36172 int ret;

36173 ctrl.buf = ctrlbuf;
36174 ctrl.len = strlen(ctrlbuf);

36175 data.buf = databuf;
36176 data.len = strlen(databuf);

36177 ret = putpmsg(fd, &ctrl, &data, 0, MSG_HIPRI);
```

36178 **APPLICATION USAGE**

36179 None.



36180 **RATIONALE**

36181 None.

36182 **FUTURE DIRECTIONS**

36183 None.

36184 **SEE ALSO**36185 *getmsg()*, *poll()*, *read()*, *write()*, the Base Definitions volume of IEEE Std 1003.1-200x, <**stropts.h**>,  
36186 Section 2.6 (on page 488)36187 **CHANGE HISTORY**

36188 First released in Issue 4, Version 2.

36189 **Issue 5**

36190 Moved from X/OPEN UNIX extension to BASE.

36191 The following text is removed from the DESCRIPTION: “The STREAM head guarantees that the  
36192 control part of a message generated by *putmsg()* is at least 64 bytes in length”.36193 **Issue 6**

36194 This function is marked as part of the XSI STREAMS Option Group.

36195 The DESCRIPTION is updated to avoid use of the term “must” for application requirements.

36196 **NAME**

36197 putpmsg — send a message on a STREAM (**STREAMS**)

36198 **SYNOPSIS**

36199 xSR #include <stropts.h>

```
36200 int putpmsg(int fildev, const struct strbuf *ctlptr,  
36201             const struct strbuf *dataptr, int band, int flags);
```

36202

36203 **DESCRIPTION**

36204 Refer to *putmsg()*.

36205 **NAME**

36206 puts — put a string on standard output

36207 **SYNOPSIS**

36208 #include &lt;stdio.h&gt;

36209 int puts(const char \*s);

36210 **DESCRIPTION**

36211 CX The functionality described on this reference page is aligned with the ISO C standard. Any  
 36212 conflict between the requirements described here and the ISO C standard is unintentional. This  
 36213 volume of IEEE Std 1003.1-200x defers to the ISO C standard.

36214 The *puts()* function shall write the string pointed to by *s*, followed by a <newline>, to the  
 36215 standard output stream *stdout*. The terminating null byte shall not be written.

36216 CX The *st\_ctime* and *st\_mtime* fields of the file shall be marked for update between the successful  
 36217 execution of *puts()* and the next successful completion of a call to *fflush()* or *fclose()* on the same  
 36218 stream or a call to *exit()* or *abort()*.

36219 **RETURN VALUE**

36220 Upon successful completion, *puts()* shall return a non-negative number. Otherwise, it shall  
 36221 CX return EOF, shall set an error indicator for the stream, and *errno* shall be set to indicate the error.

36222 **ERRORS**36223 Refer to *fputc()*.36224 **EXAMPLES**36225 **Printing to Standard Output**

36226 The following example gets the current time, converts it to a string using *localtime()* and  
 36227 *asctime()*, and prints it to standard output using *puts()*. It then prints the number of minutes to  
 36228 an event for which it is waiting.

```

36229 #include <time.h>
36230 #include <stdio.h>
36231 ...
36232 time_t now;
36233 int minutes_to_event;
36234 ...
36235 time(&now);
36236 printf("The time is ");
36237 puts(asctime(localtime(&now)));
36238 printf("There are %d minutes to the event.\n",
36239     minutes_to_event);
36240 ...

```

36241 **APPLICATION USAGE**36242 The *puts()* function appends a <newline>, while *fputs()* does not.36243 **RATIONALE**

36244 None.

36245 **FUTURE DIRECTIONS**

36246 None.

36247 **SEE ALSO**

36248 *fopen()*, *fputs()*, *putc()*, the Base Definitions volume of IEEE Std 1003.1-200x, <**stdio.h**>

36249 **CHANGE HISTORY**

36250 First released in Issue 1. Derived from Issue 1 of the SVID.

36251 **Issue 6**

36252 Extensions beyond the ISO C standard are now marked.

36253 **NAME**

36254 pututxline — put an entry into user accounting database

36255 **SYNOPSIS**36256 XSI `#include <utmpx.h>`36257 `struct utmpx *pututxline(const struct utmpx *utmpx);`

36258

36259 **DESCRIPTION**36260 Refer to *endutxent()*.

36261 **NAME**

36262 putwc — put a wide character on a stream

36263 **SYNOPSIS**

36264 #include &lt;stdio.h&gt;

36265 #include &lt;wchar.h&gt;

36266 wint\_t putwc(wchar\_t *wc*, FILE \**stream*);36267 **DESCRIPTION**

36268 cx The functionality described on this reference page is aligned with the ISO C standard. Any  
36269 conflict between the requirements described here and the ISO C standard is unintentional. This  
36270 volume of IEEE Std 1003.1-200x defers to the ISO C standard.

36271 The *putwc()* function shall be equivalent to *fputwc()*, except that if it is implemented as a macro  
36272 it may evaluate *stream* more than once, so the argument should never be an expression with side  
36273 effects.

36274 **RETURN VALUE**36275 Refer to *fputwc()*.36276 **ERRORS**36277 Refer to *fputwc()*.36278 **EXAMPLES**

36279 None.

36280 **APPLICATION USAGE**

36281 Since it may be implemented as a macro, *putwc()* may treat a *stream* argument with side effects  
36282 incorrectly. In particular, *putwc(wc,\*f++)* need not work correctly. Therefore, use of this function  
36283 is not recommended; *fputwc()* should be used instead.

36284 **RATIONALE**

36285 None.

36286 **FUTURE DIRECTIONS**

36287 None.

36288 **SEE ALSO**36289 *fputwc()*, the Base Definitions volume of IEEE Std 1003.1-200x, <stdio.h>, <wchar.h>36290 **CHANGE HISTORY**

36291 First released as a World-wide Portability Interface in Issue 4.

36292 **Issue 5**

36293 Aligned with ISO/IEC 9899:1990/Amendment 1:1995 (E). Specifically, the type of argument *wc*  
36294 is changed from **wint\_t** to **wchar\_t**.

36295 The Optional Header (OH) marking is removed from &lt;stdio.h&gt;.

36296 **NAME**

36297 putwchar — put a wide character on stdout stream

36298 **SYNOPSIS**

36299 #include <wchar.h>

36300 wint\_t putwchar(wchar\_t wc);

36301 **DESCRIPTION**

36302 cx The functionality described on this reference page is aligned with the ISO C standard. Any  
36303 conflict between the requirements described here and the ISO C standard is unintentional. This  
36304 volume of IEEE Std 1003.1-200x defers to the ISO C standard.

36305 The function call *putwchar(wc)* shall be equivalent to *putwc(wc,stdout)*.

36306 **RETURN VALUE**

36307 Refer to *fputwc()*.

36308 **ERRORS**

36309 Refer to *fputwc()*.

36310 **EXAMPLES**

36311 None.

36312 **APPLICATION USAGE**

36313 None.

36314 **RATIONALE**

36315 None.

36316 **FUTURE DIRECTIONS**

36317 None.

36318 **SEE ALSO**

36319 *fputwc()*, *putwc()*, the Base Definitions volume of IEEE Std 1003.1-200x, <wchar.h>

36320 **CHANGE HISTORY**

36321 First released in Issue 4.

36322 **Issue 5**

36323 Aligned with ISO/IEC 9899:1990/Amendment 1:1995 (E). Specifically, the type of argument *wc*  
36324 is changed from **wint\_t** to **wchar\_t**.

36325 **NAME**

36326           pwrite — write on a file

36327 **SYNOPSIS**

36328           #include <unistd.h>

```
36329 xSI       ssize_t pwrite(int fildes, const void *buf, size_t nbyte,  
36330               off_t offset);
```

36331

36332 **DESCRIPTION**

36333           Refer to *write()*.



36334 **NAME**

36335 qsort — sort a table of data

36336 **SYNOPSIS**

36337 #include &lt;stdlib.h&gt;

36338 void qsort(void \*base, size\_t nel, size\_t width,  
36339 int (\*compar)(const void \*, const void \*));36340 **DESCRIPTION**36341 cx The functionality described on this reference page is aligned with the ISO C standard. Any  
36342 conflict between the requirements described here and the ISO C standard is unintentional. This  
36343 volume of IEEE Std 1003.1-200x defers to the ISO C standard.36344 The *qsort()* function shall sort an array of *nel* objects, the initial element of which is pointed to by  
36345 *base*. The size of each object, in bytes, is specified by the *width* argument.36346 The contents of the array shall be sorted in ascending order according to a comparison function.  
36347 The *compar* argument is a pointer to the comparison function, which is called with two  
36348 arguments that point to the elements being compared. The application shall ensure that the  
36349 function returns an integer less than, equal to, or greater than 0, if the first argument is  
36350 considered respectively less than, equal to, or greater than the second. If two members compare  
36351 as equal, their order in the sorted array is unspecified.36352 **RETURN VALUE**36353 The *qsort()* function shall not return a value.36354 **ERRORS**

36355 No errors are defined.

36356 **EXAMPLES**

36357 None.

36358 **APPLICATION USAGE**36359 The comparison function need not compare every byte, so arbitrary data may be contained in  
36360 the elements in addition to the values being compared.36361 **RATIONALE**

36362 None.

36363 **FUTURE DIRECTIONS**

36364 None.

36365 **SEE ALSO**

36366 The Base Definitions volume of IEEE Std 1003.1-200x, &lt;stdlib.h&gt;

36367 **CHANGE HISTORY**

36368 First released in Issue 1. Derived from Issue 1 of the SVID.

36369 **Issue 6**

36370 The DESCRIPTION is updated to avoid use of the term “must” for application requirements.

36371 **NAME**

36372 raise — send a signal to the executing process

36373 **SYNOPSIS**

36374 #include <signal.h>

36375 int raise(int sig);

36376 **DESCRIPTION**

36377 CX The functionality described on this reference page is aligned with the ISO C standard. Any  
 36378 conflict between the requirements described here and the ISO C standard is unintentional. This  
 36379 volume of IEEE Std 1003.1-200x defers to the ISO C standard.

36380 CX The *raise()* function shall send the signal *sig* to the executing thread or process. If a signal  
 36381 handler is called, the *raise()* function shall not return until after the signal handler does.

36382 THR If the implementation supports the Threads option, the effect of the *raise()* function shall be  
 36383 equivalent to calling:

36384 pthread\_kill(pthread\_self(), sig);

36385

36386 CX Otherwise, the effect of the *raise()* function shall be equivalent to calling:

36387 kill(getpid(), sig);

36388

36389 **RETURN VALUE**

36390 CX Upon successful completion, 0 shall be returned. Otherwise, a non-zero value shall be returned  
 36391 and *errno* shall be set to indicate the error.

36392 **ERRORS**

36393 The *raise()* function shall fail if:

36394 CX [EINVAL] The value of the *sig* argument is an invalid signal number.

36395 **EXAMPLES**

36396 None.

36397 **APPLICATION USAGE**

36398 None.

36399 **RATIONALE**

36400 The term “thread” is an extension to the ISO C standard.

36401 **FUTURE DIRECTIONS**

36402 None.

36403 **SEE ALSO**

36404 *kill()*, *sigaction()*, the Base Definitions volume of IEEE Std 1003.1-200x, <signal.h>,  
 36405 <sys/types.h>

36406 **CHANGE HISTORY**

36407 First released in Issue 4. Derived from the ANSI C standard.

36408 **Issue 5**

36409 The DESCRIPTION is updated for alignment with the POSIX Threads Extension.

36410 **Issue 6**

36411 Extensions beyond the ISO C standard are now marked.

36412 The following new requirements on POSIX implementations derive from alignment with the  
36413 Single UNIX Specification:

- 36414 • In the RETURN VALUE section, the requirement to set *errno* on error is added.
- 36415 • The [EINVAL] error condition is added.

## 36416 NAME

36417 rand, rand\_r, srand — pseudo-random number generator

## 36418 SYNOPSIS

36419 #include &lt;stdlib.h&gt;

36420 int rand(void);

36421 TSF int rand\_r(unsigned \*seed);

36422 void srand(unsigned seed);

## 36423 DESCRIPTION

36424 CX The functionality described on this reference page is aligned with the ISO C standard. Any  
 36425 conflict between the requirements described here and the ISO C standard is unintentional. This  
 36426 volume of IEEE Std 1003.1-200x defers to the ISO C standard.

36427 The *rand()* function shall compute a sequence of pseudo-random integers in the range 0 to  
 36428 XSI {RAND\_MAX} with a period of at least  $2^{32}$ .

36429 CX The *rand()* function need not be reentrant. A function that is not required to be reentrant is not  
 36430 required to be thread-safe.

36431 TSF The *rand\_r()* function shall compute a sequence of pseudo-random integers in the range 0 to  
 36432 {RAND\_MAX}. (The value of the {RAND\_MAX} macro shall be at least 32 767.)

36433 If *rand\_r()* is called with the same initial value for the object pointed to by *seed* and that object is  
 36434 not modified between successive returns and calls to *rand\_r()*, the same sequence shall be  
 36435 generated.

36436 The *srand()* function uses the argument as a seed for a new sequence of pseudo-random  
 36437 numbers to be returned by subsequent calls to *rand()*. If *srand()* is then called with the same  
 36438 seed value, the sequence of pseudo-random numbers shall be repeated. If *rand()* is called before  
 36439 any calls to *srand()* are made, the same sequence shall be generated as when *srand()* is first  
 36440 called with a seed value of 1.

36441 The implementation shall behave as if no function defined in this volume of  
 36442 IEEE Std 1003.1-200x calls *rand()* or *srand()*.

## 36443 RETURN VALUE

36444 The *rand()* function shall return the next pseudo-random number in the sequence.36445 TSF The *rand\_r()* function shall return a pseudo-random integer.36446 The *srand()* function shall not return a value.

## 36447 ERRORS

36448 No errors are defined.

## 36449 EXAMPLES

36450 **Generating a Pseudo-Random Number Sequence**

36451 The following example demonstrates how to generate a sequence of pseudo-random numbers.

36452 #include &lt;stdio.h&gt;

36453 #include &lt;stdlib.h&gt;

36454 ...

36455 long count, i;

36456 char \*keyst;

36457 int elementlen, len;

36458 char c;

```

36459     ...
36460     /* Initial random number generator. */
36461     srand(1);

36462     /* Create keys using only lower case characters */
36463     len = 0;
36464     for (i=0; i<count; i++) {
36465         while (len < elementlen) {
36466             c = (char) (rand() % 128);
36467             if (islower(c))
36468                 keystr[len++] = c;
36469         }

36470         keystr[len] = '\0';
36471         printf("%s Element%0*ld\n", keystr, elementlen, i);
36472         len = 0;
36473     }

```

### 36474 **Generating the Same Sequence on Different Machines**

36475 The following code defines a pair of functions that could be incorporated into applications  
 36476 wishing to ensure that the same sequence of numbers is generated across different machines.

```

36477     static unsigned long next = 1;
36478     int myrand(void) /* RAND_MAX assumed to be 32767. */
36479     {
36480         next = next * 1103515245 + 12345;
36481         return((unsigned)(next/65536) % 32768);
36482     }

36483     void mysrand(unsigned seed)
36484     {
36485         next = seed;
36486     }

```

### 36487 **APPLICATION USAGE**

36488 The *drand48()* function provides a much more elaborate random number generator.

### 36489 **RATIONALE**

36490 The ISO C standard *rand()* and *srand()* functions allow per-process pseudo-random streams  
 36491 shared by all threads. Those two functions need not change, but there has to be mutual-  
 36492 exclusion that prevents interference between two threads concurrently accessing the random  
 36493 number generator.

36494 With regard to *rand()*, there are two different behaviors that may be wanted in a multi-threaded  
 36495 program:

- 36496 1. A single per-process sequence of pseudo-random numbers that is shared by all threads  
 36497 that call *rand()*
- 36498 2. A different sequence of pseudo-random numbers for each thread that calls *rand()*

36499 This is provided by the modified thread-safe function based on whether the seed value is global  
 36500 to the entire process or local to each thread.

36501 This does not address the known deficiencies of the *rand()* function implementations, which  
 36502 have been approached by maintaining more state. In effect, this specifies new thread-safe forms  
 36503 of a deficient function.

36504 **FUTURE DIRECTIONS**

36505           None.

36506 **SEE ALSO**36507           *drand48()*, the Base Definitions volume of IEEE Std 1003.1-200x, <**stdlib.h**>36508 **CHANGE HISTORY**

36509           First released in Issue 1. Derived from Issue 1 of the SVID.

36510 **Issue 5**36511           The *rand\_r()* function is included for alignment with the POSIX Threads Extension.36512           A note indicating that the *rand()* function need not be reentrant is added to the DESCRIPTION.36513 **Issue 6**

36514           Extensions beyond the ISO C standard are now marked.

36515           The *rand\_r()* function is marked as part of the Thread-Safe Functions option.

36516 **NAME**

36517 random — generate pseudo-random number

36518 **SYNOPSIS**36519 xSI `#include <stdlib.h>`36520 `long random(void);`

36521

36522 **DESCRIPTION**36523 Refer to *initstate()*.

36524 **NAME**

36525 pread, read — read from a file

36526 **SYNOPSIS**

36527 #include &lt;unistd.h&gt;

36528 XSI `ssize_t pread(int fildev, void *buf, size_t nbyte, off_t offset);`36529 `ssize_t read(int fildev, void *buf, size_t nbyte);`36530 **DESCRIPTION**

36531 The `read()` function shall attempt to read *nbyte* bytes from the file associated with the open file  
 36532 descriptor, *fildev*, into the buffer pointed to by *buf*. The behavior of multiple concurrent reads on  
 36533 the same pipe, FIFO, or terminal device is unspecified.

36534 Before any action described below is taken, and if *nbyte* is zero, the `read()` function may detect  
 36535 and return errors as described below. In the absence of errors, or if error detection is not  
 36536 performed, the `read()` function shall return zero and have no other results.

36537 On files that support seeking (for example, a regular file), the `read()` shall start at a position in  
 36538 the file given by the file offset associated with *fildev*. The file offset shall be incremented by the  
 36539 number of bytes actually read.

36540 Files that do not support seeking—for example, terminals—always read from the current  
 36541 position. The value of a file offset associated with such a file is undefined.

36542 No data transfer shall occur past the current end-of-file. If the starting position is at or after the  
 36543 end-of-file, 0 shall be returned. If the file refers to a device special file, the result of subsequent  
 36544 `read()` requests is implementation-defined.

36545 If the value of *nbyte* is greater than {SSIZE\_MAX}, the result is implementation-defined.

36546 When attempting to read from an empty pipe or FIFO:

- 36547 • If no process has the pipe open for writing, `read()` shall return 0 to indicate end-of-file.
- 36548 • If some process has the pipe open for writing and O\_NONBLOCK is set, `read()` shall return  
 36549 -1 and set *errno* to [EAGAIN].
- 36550 • If some process has the pipe open for writing and O\_NONBLOCK is clear, `read()` shall block  
 36551 the calling thread until some data is written or the pipe is closed by all processes that had the  
 36552 pipe open for writing.

36553 When attempting to read a file (other than a pipe or FIFO) that supports non-blocking reads and  
 36554 has no data currently available:

- 36555 • If O\_NONBLOCK is set, `read()` shall return -1 and set *errno* to [EAGAIN].
- 36556 • If O\_NONBLOCK is clear, `read()` shall block the calling thread until some data becomes  
 36557 available.
- 36558 • The use of the O\_NONBLOCK flag has no effect if there is some data available.

36559 The `read()` function reads data previously written to a file. If any portion of a regular file prior to  
 36560 the end-of-file has not been written, `read()` shall return bytes with value 0. For example, `lseek()`  
 36561 allows the file offset to be set beyond the end of existing data in the file. If data is later written at  
 36562 this point, subsequent reads in the gap between the previous end of data and the newly written  
 36563 data shall return bytes with value 0 until data is written into the gap.

36564 Upon successful completion, where *nbyte* is greater than 0, `read()` shall mark for update the  
 36565 *st\_atime* field of the file, and shall return the number of bytes read. This number shall never be  
 36566 greater than *nbyte*. The value returned may be less than *nbyte* if the number of bytes left in the



36567 file is less than *nbyte*, if the *read()* request was interrupted by a signal, or if the file is a pipe or  
 36568 FIFO or special file and has fewer than *nbyte* bytes immediately available for reading. For  
 36569 example, a *read()* from a file associated with a terminal may return one typed line of data.

36570 If a *read()* is interrupted by a signal before it reads any data, it shall return  $-1$  with *errno* set to  
 36571 [EINTR].

36572 If a *read()* is interrupted by a signal after it has successfully read some data, it shall return the  
 36573 number of bytes read.

36574 For regular files, no data transfer shall occur past the offset maximum established in the open  
 36575 file description associated with *fdes*.

36576 If *fdes* refers to a socket, *read()* shall be equivalent to *recv()* with no flags set. |

36577 SIO If the O\_DSYNC and O\_RSYNC bits have been set, read I/O operations on the file descriptor |  
 36578 shall complete as defined by synchronized I/O data integrity completion. If the O\_SYNC and |  
 36579 O\_RSYNC bits have been set, read I/O operations on the file descriptor shall complete as |  
 36580 defined by synchronized I/O file integrity completion. |

36581 SHM If *fdes* refers to a shared memory object, the result of the *read()* function is unspecified. |

36582 TYM If *fdes* refers to a typed memory object, the result of the *read()* function is unspecified. |

36583 XSR A *read()* from a STREAMS file can read data in three different modes: *byte-stream* mode, |  
 36584 *message-nondiscard* mode, and *message-discard* mode. The default shall be byte-stream mode. This |  
 36585 can be changed using the I\_SRDOPT *ioctl()* request, and can be tested with the I\_GRDOPT |  
 36586 *ioctl()*. In byte-stream mode, *read()* shall retrieve data from the STREAM until as many bytes as |  
 36587 were requested are transferred, or until there is no more data to be retrieved. Byte-stream mode |  
 36588 ignores message boundaries. |

36589 In STREAMS message-nondiscard mode, *read()* shall retrieve data until as many bytes as were  
 36590 requested are transferred, or until a message boundary is reached. If *read()* does not retrieve all  
 36591 the data in a message, the remaining data shall be left on the STREAM, and can be retrieved by  
 36592 the next *read()* call. Message-discard mode also retrieves data until as many bytes as were  
 36593 requested are transferred, or a message boundary is reached. However, unread data remaining  
 36594 in a message after the *read()* returns shall be discarded, and shall not be available for a  
 36595 subsequent *read()*, *getmsg()*, or *getpmsg()* call. |

36596 How *read()* handles zero-byte STREAMS messages is determined by the current read mode  
 36597 setting. In byte-stream mode, *read()* shall accept data until it has read *nbyte* bytes, or until there  
 36598 is no more data to read, or until a zero-byte message block is encountered. The *read()* function  
 36599 shall then return the number of bytes read, and place the zero-byte message back on the  
 36600 STREAM to be retrieved by the next *read()*, *getmsg()*, or *getpmsg()*. In message-nondiscard mode |  
 36601 or message-discard mode, a zero-byte message shall return 0 and the message shall be removed  
 36602 from the STREAM. When a zero-byte message is read as the first message on a STREAM, the  
 36603 message shall be removed from the STREAM and 0 shall be returned, regardless of the read  
 36604 mode. |

36605 A *read()* from a STREAMS file shall return the data in the message at the front of the STREAM  
 36606 head read queue, regardless of the priority band of the message. |

36607 By default, STREAMS are in control-normal mode, in which a *read()* from a STREAMS file can  
 36608 only process messages that contain a data part but do not contain a control part. The *read()* shall  
 36609 fail if a message containing a control part is encountered at the STREAM head. This default  
 36610 action can be changed by placing the STREAM in either control-data mode or control-discard  
 36611 mode with the I\_SRDOPT *ioctl()* command. In control-data mode, *read()* shall convert any  
 36612 control part to data and pass it to the application before passing any data part originally present

36613 in the same message. In control-discard mode, *read()* shall discard message control parts but  
36614 return to the process any data part in the message.

36615 In addition, *read()* shall fail if the STREAM head had processed an asynchronous error before the  
36616 call. In this case, the value of *errno* shall not reflect the result of *read()*, but reflects the prior error.  
36617 If a hangup occurs on the STREAM being read, *read()* shall continue to operate normally until  
36618 the STREAM head read queue is empty. Thereafter, it shall return 0.

36619 XSI The *pread()* function shall be equivalent to *read()*, except that it shall read from a given position  
36620 in the file without changing the file pointer. The first three arguments to *pread()* are the same as  
36621 *read()* with the addition of a fourth argument offset for the desired position inside the file. An  
36622 attempt to perform a *pread()* on a file that is incapable of seeking shall result in an error.

### 36623 RETURN VALUE

36624 XSI Upon successful completion, *read()* and *pread()* shall return a non-negative integer indicating the  
36625 number of bytes actually read. Otherwise, the functions shall return -1 and set *errno* to indicate  
36626 the error.

### 36627 ERRORS

36628 XSI The *read()* and *pread()* functions shall fail if:

36629 [EAGAIN] The O\_NONBLOCK flag is set for the file descriptor and the process would be  
36630 delayed.

36631 [EBADF] The *fildev* argument is not a valid file descriptor open for reading.

36632 XSR [EBADMSG] The file is a STREAM file that is set to control-normal mode and the message  
36633 waiting to be read includes a control part.

36634 [EINTR] The read operation was terminated due to the receipt of a signal, and no data  
36635 was transferred.

36636 XSR [EINVAL] The STREAM or multiplexer referenced by *fildev* is linked (directly or  
36637 indirectly) downstream from a multiplexer.

36638 [EIO] The process is a member of a background process attempting to read from its  
36639 controlling terminal, the process is ignoring or blocking the SIGTTIN signal,  
36640 or the process group is orphaned. This error may also be generated for  
36641 implementation-defined reasons.

36642 XSI [EISDIR] The *fildev* argument refers to a directory and the implementation does not  
36643 allow the directory to be read using *read()* or *pread()*. The *readdir()* function  
36644 should be used instead.

36645 [E\_OVERFLOW] The file is a regular file, *nbyte* is greater than 0, the starting position is before  
36646 the end-of-file, and the starting position is greater than or equal to the offset  
36647 maximum established in the open file description associated with *fildev*.

36648 The *read()* function shall fail if:

36649 [EAGAIN] or [EWOULDBLOCK]

36650 The file descriptor is for a socket, is marked O\_NONBLOCK, and no data is  
36651 waiting to be received.

36652 [ECONNRESET] A read was attempted on a socket and the connection was forcibly closed by  
36653 its peer.

36654 [ENOTCONN] A read was attempted on a socket that is not connected.

36655 [ETIMEDOUT] A read was attempted on a socket and a transmission timeout occurred.

36656 XSI	The <code>read()</code> and <code>pread()</code> functions may fail if:	
36657	[EIO] A physical I/O error has occurred.	
36658	[ENOBUFS] Insufficient resources were available in the system to perform the operation.	
36659	[ENOMEM] Insufficient memory was available to fulfill the request.	
36660	[ENXIO] A request was made of a nonexistent device, or the request was outside the	
36661	capabilities of the device.	
36662	The <code>pread()</code> function shall fail, and the file pointer shall remain unchanged, if:	
36663 XSI	[EINVAL] The <code>offset</code> argument is invalid. The value is negative.	
36664 XSI	[EOVERFLOW] The file is a regular file and an attempt was made to read at or beyond the	
36665	offset maximum associated with the file.	
36666 XSI	[ENXIO] A request was outside the capabilities of the device.	
36667 XSI	[ESPIPE] <code>fildev</code> is associated with a pipe or FIFO.	

### 36668 EXAMPLES

#### 36669 Reading Data into a Buffer

36670 The following example reads data from the file associated with the file descriptor `fd` into the  
 36671 buffer pointed to by `buf`.

```

36672 #include <sys/types.h>
36673 #include <unistd.h>
36674 ...
36675 char buf[20];
36676 size_t nbytes;
36677 ssize_t bytes_read;
36678 int fd;
36679 ...
36680 nbytes = sizeof(buf);
36681 bytes_read = read(fd, buf, nbytes);
36682 ...

```

### 36683 APPLICATION USAGE

36684 None.

### 36685 RATIONALE

36686 This volume of IEEE Std 1003.1-200x does not specify the value of the file offset after an error is  
 36687 returned; there are too many cases. For programming errors, such as [EBADF], the concept is  
 36688 meaningless since no file is involved. For errors that are detected immediately, such as  
 36689 [EAGAIN], clearly the pointer should not change. After an interrupt or hardware error, however,  
 36690 an updated value would be very useful and is the behavior of many implementations.

36691 Note that a `read()` of zero bytes does not modify `st_atime`. A `read()` that requests more than zero  
 36692 bytes, but returns zero, shall modify `st_atime`.

36693 Implementations are allowed, but not required, to perform error checking for `read()` requests of  
 36694 zero bytes.

36695 **Input and Output**

36696 The use of I/O with large byte counts has always presented problems. Ideas such as *lread()* and  
36697 *lwrite()* (using and returning **longs**) were considered at one time. The current solution is to use  
36698 abstract types on the ISO C standard function to *read()* and *write()*. The abstract types can be  
36699 declared so that existing functions work, but can also be declared so that larger types can be  
36700 represented in future implementations. It is presumed that whatever constraints limit the  
36701 maximum range of **size\_t** also limit portable I/O requests to the same range. This volume of  
36702 IEEE Std 1003.1-200x also limits the range further by requiring that the byte count be limited so  
36703 that a signed return value remains meaningful. Since the return type is also a (signed) abstract  
36704 type, the byte count can be defined by the implementation to be larger than an **int** can hold.

36705 The standard developers considered adding atomicity requirements to a pipe or FIFO, but  
36706 recognized that due to the nature of pipes and FIFOs there could be no guarantee of atomicity of  
36707 reads of {PIPE\_BUF} or any other size that would be an aid to applications portability.

36708 This volume of IEEE Std 1003.1-200x requires that no action be taken for *read()* or *write()* when  
36709 *nbyte* is zero. This is not intended to take precedence over detection of errors (such as invalid  
36710 buffer pointers or file descriptors). This is consistent with the rest of this volume of  
36711 IEEE Std 1003.1-200x, but the phrasing here could be misread to require detection of the zero  
36712 case before any other errors. A value of zero is to be considered a correct value, for which the  
36713 semantics are a no-op.

36714 I/O is intended to be atomic to ordinary files and pipes and FIFOs. Atomic means that all the  
36715 bytes from a single operation that started out together end up together, without interleaving  
36716 from other I/O operations. It is a known attribute of terminals that this is not honored, and  
36717 terminals are explicitly (and implicitly permanently) excepted, making the behavior unspecified.  
36718 The behavior for other device types is also left unspecified, but the wording is intended to imply  
36719 that future standards might choose to specify atomicity (or not).

36720 There were recommendations to add format parameters to *read()* and *write()* in order to handle  
36721 networked transfers among heterogeneous file system and base hardware types. Such a facility  
36722 may be required for support by the OSI presentation of layer services. However, it was  
36723 determined that this should correspond with similar C-language facilities, and that is beyond the  
36724 scope of this volume of IEEE Std 1003.1-200x. The concept was suggested to the developers of  
36725 the ISO C standard for their consideration as a possible area for future work.

36726 In 4.3 BSD, a *read()* or *write()* that is interrupted by a signal before transferring any data does not  
36727 by default return an [EINTR] error, but is restarted. In 4.2 BSD, 4.3 BSD, and the Eighth Edition,  
36728 there is an additional function, *select()*, whose purpose is to pause until specified activity (data  
36729 to read, space to write, and so on) is detected on specified file descriptors. It is common in  
36730 applications written for those systems for *select()* to be used before *read()* in situations (such as  
36731 keyboard input) where interruption of I/O due to a signal is desired.

36732 The issue of which files or file types are interruptible is considered an implementation design  
36733 issue. This is often affected primarily by hardware and reliability issues.

36734 There are no references to actions taken following an “unrecoverable error”. It is considered  
36735 beyond the scope of this volume of IEEE Std 1003.1-200x to describe what happens in the case of  
36736 hardware errors.

36737 Previous versions of IEEE Std 1003.1-200x allowed two very different behaviors with regard to  
36738 the handling of interrupts. In order to minimize the resulting confusion, it was decided that  
36739 IEEE Std 1003.1-200x should support only one of these behaviors. Historical practice on AT&T-  
36740 derived systems was to have *read()* and *write()* return **-1** and set *errno* to [EINTR] when  
36741 interrupted after some, but not all, of the data requested had been transferred. However, the U.S.  
36742 Department of Commerce FIPS 151-1 and FIPS 151-2 require the historical BSD behavior, in

36743 which *read()* and *write()* return the number of bytes actually transferred before the interrupt. If  
 36744  $-1$  is returned when any data is transferred, it is difficult to recover from the error on a seekable  
 36745 device and impossible on a non-seekable device. Most new implementations support this  
 36746 behavior. The behavior required by IEEE Std 1003.1-200x is to return the number of bytes  
 36747 transferred.

36748 IEEE Std 1003.1-200x does not specify when an implementation that buffers *read()*s actually  
 36749 moves the data into the user-supplied buffer, so an implementation may chose to do this at the  
 36750 latest possible moment. Therefore, an interrupt arriving earlier may not cause *read()* to return a  
 36751 partial byte count, but rather to return  $-1$  and set *errno* to [EINTR].

36752 Consideration was also given to combining the two previous options, and setting *errno* to  
 36753 [EINTR] while returning a short count. However, not only is there no existing practice that  
 36754 implements this, it is also contradictory to the idea that when *errno* is set, the function  
 36755 responsible shall return  $-1$ .

#### 36756 FUTURE DIRECTIONS

36757 None.

#### 36758 SEE ALSO

36759 *fcntl()*, *ioctl()*, *lseek()*, *open()*, *pipe()*, *readv()*, the Base Definitions volume of |  
 36760 IEEE Std 1003.1-200x, <stropts.h>, <sys/uio.h>, <unistd.h>, the Base Definitions volume of  
 36761 IEEE Std 1003.1-200x, Chapter 11, General Terminal Interface

#### 36762 CHANGE HISTORY

36763 First released in Issue 1. Derived from Issue 1 of the SVID.

#### 36764 Issue 5

36765 The DESCRIPTION is updated for alignment with the POSIX Realtime Extension and the POSIX  
 36766 Threads Extension.

36767 Large File Summit extensions are added.

36768 The *pread()* function is added.

#### 36769 Issue 6

36770 The DESCRIPTION and ERRORS sections are updated so that references to STREAMS are  
 36771 marked as part of the XSI STREAMS Option Group.

36772 The following new requirements on POSIX implementations derive from alignment with the  
 36773 Single UNIX Specification:

36774 • The DESCRIPTION now states that if *read()* is interrupted by a signal after it has successfully  
 36775 read some data, it returns the number of bytes read. In Issue 3, it was optional whether *read()*  
 36776 returned the number of bytes read, or whether it returned  $-1$  with *errno* set to [EINTR]. This  
 36777 is a FIPS requirement.

36778 • In the DESCRIPTION, text is added to indicate that for regular files, no data transfer occurs  
 36779 past the offset maximum established in the open file description associated with *files*. This  
 36780 change is to support large files.

36781 • The [EOVERFLOW] mandatory error condition is added.

36782 • The [ENXIO] optional error condition is added.

36783 Text referring to sockets is added to the DESCRIPTION.

36784 The following changes were made to align with the IEEE P1003.1a draft standard:

36785 • The effect of reading zero bytes is clarified.

36786 The DESCRIPTION is updated for alignment with IEEE Std 1003.1j-2000 by specifying that  
36787 *read()* results are unspecified for typed memory objects.

36788 New RATIONALE is added to explain the atomicity requirements for input and output  
36789 operations.

36790 The following error conditions are added for operations on sockets: [EAGAIN],  
36791 [ECONNRESET], [ENOTCONN], and [ETIMEDOUT].

36792 The [EIO] error is changed to “may fail”.

36793 The following error conditions are added for operations on sockets: [ENOBUFFS] and  
36794 [ENOMEM]. |

36795 The *readv()* function is split out into a separate reference page. |

## 36796 NAME

36797 readdir, readdir\_r — read directory

## 36798 SYNOPSIS

36799 #include &lt;dirent.h&gt;

36800 struct dirent \*readdir(DIR \*dirp);

36801 TSF int readdir\_r(DIR \*restrict dirp, struct dirent \*restrict entry,

36802 struct dirent \*\*restrict result);

36803

## 36804 DESCRIPTION

36805 The type **DIR**, which is defined in the <**dirent.h**> header, represents a *directory stream*, which is  
 36806 an ordered sequence of all the directory entries in a particular directory. Directory entries  
 36807 represent files; files may be removed from a directory or added to a directory asynchronously to  
 36808 the operation of *readdir()*.

36809 The *readdir()* function shall return a pointer to a structure representing the directory entry at the  
 36810 current position in the directory stream specified by the argument *dirp*, and position the  
 36811 directory stream at the next entry. It shall return a null pointer upon reaching the end of the  
 36812 directory stream. The structure **dirent** defined in the <**dirent.h**> header describes a directory  
 36813 entry.

36814 The *readdir()* function shall not return directory entries containing empty names. If entries for  
 36815 dot or dot-dot exist, one entry shall be returned for dot and one entry shall be returned for dot-  
 36816 dot; otherwise, they shall not be returned.

36817 The pointer returned by *readdir()* points to data which may be overwritten by another call to  
 36818 *readdir()* on the same directory stream. This data is not overwritten by another call to *readdir()*  
 36819 on a different directory stream.

36820 If a file is removed from or added to the directory after the most recent call to *opendir()* or  
 36821 *rewinddir()*, whether a subsequent call to *readdir()* returns an entry for that file is unspecified.

36822 The *readdir()* function may buffer several directory entries per actual read operation; *readdir()*  
 36823 shall mark for update the *st\_atime* field of the directory each time the directory is actually read.

36824 After a call to *fork()*, either the parent or child (but not both) may continue processing the  
 36825 XSI directory stream using *readdir()*, *rewinddir()*, or *seekdir()*. If both the parent and child processes  
 36826 use these functions, the result is undefined.

36827 If the entry names a symbolic link, the value of the *d\_ino* member is unspecified.

36828 The *readdir()* function need not be reentrant. A function that is not required to be reentrant is not  
 36829 required to be thread-safe.

36830 TSF The *readdir\_r()* function shall initialize the **dirent** structure referenced by *entry* to represent the  
 36831 directory entry at the current position in the directory stream referred to by *dirp*, store a pointer  
 36832 to this structure at the location referenced by *result*, and position the directory stream at the next  
 36833 entry.

36834 The storage pointed to by *entry* shall be large enough for a **dirent** with an array of **char** *d\_name*  
 36835 members containing at least {NAME\_MAX} plus one elements.

36836 Upon successful return, the pointer returned at *\*result* shall have the same value as the argument  
 36837 *entry*. Upon reaching the end of the directory stream, this pointer shall have the value NULL.

36838 The *readdir\_r()* function shall not return directory entries containing empty names.

36839 If a file is removed from or added to the directory after the most recent call to *opendir()* or  
 36840 *rewinddir()*, whether a subsequent call to *readdir\_r()* returns an entry for that file is unspecified.

36841 The *readdir\_r()* function may buffer several directory entries per actual read operation; the  
 36842 *readdir\_r()* function shall mark for update the *st\_atime* field of the directory each time the  
 36843 directory is actually read.

36844 Applications wishing to check for error situations should set *errno* to 0 before calling *readdir()*. If  
 36845 *errno* is set to non-zero on return, an error occurred.

#### 36846 RETURN VALUE

36847 Upon successful completion, *readdir()* shall return a pointer to an object of type **struct dirent**.  
 36848 When an error is encountered, a null pointer shall be returned and *errno* shall be set to indicate  
 36849 the error. When the end of the directory is encountered, a null pointer shall be returned and *errno*  
 36850 is not changed.

36851 TSF If successful, the *readdir\_r()* function shall return zero; otherwise, an error number shall be  
 36852 returned to indicate the error.

#### 36853 ERRORS

36854 The *readdir()* function shall fail if:

36855 [EOverflow] One of the values in the structure to be returned cannot be represented  
 36856 correctly.

36857 The *readdir()* function may fail if:

36858 [EBADF] The *dirp* argument does not refer to an open directory stream.

36859 [ENOENT] The current position of the directory stream is invalid.

36860 The *readdir\_r()* function may fail if:

36861 [EBADF] The *dirp* argument does not refer to an open directory stream.

#### 36862 EXAMPLES

36863 The following sample code searches the current directory for the entry *name*:

```
36864 dirp = opendir(".");
36865 while (dirp) {
36866     errno = 0;
36867     if ((dp = readdir(dirp)) != NULL) {
36868         if (strcmp(dp->d_name, name) == 0) {
36869             closedir(dirp);
36870             return FOUND;
36871         }
36872     } else {
36873         if (errno == 0) {
36874             closedir(dirp);
36875             return NOT_FOUND;
36876         }
36877         closedir(dirp);
36878         return READ_ERROR;
36879     }
36880 }
36881 return OPEN_ERROR;
```



36882 **APPLICATION USAGE**

36883       The *readdir()* function should be used in conjunction with *opendir()*, *closedir()*, and *rewinddir()* to  
36884       examine the contents of the directory.

36885       The *readdir\_r()* function is thread-safe and shall return values in a user-supplied buffer instead  
36886       of possibly using a static data area that may be overwritten by each call.

36887 **RATIONALE**

36888       The returned value of *readdir()* merely *represents* a directory entry. No equivalence should be  
36889       inferred.

36890       Historical implementations of *readdir()* obtain multiple directory entries on a single read  
36891       operation, which permits subsequent *readdir()* operations to operate from the buffered  
36892       information. Any wording that required each successful *readdir()* operation to mark the  
36893       directory *st\_atime* field for update would militate against the historical performance-oriented  
36894       implementations.

36895       Since *readdir()* returns NULL when it detects an error and when the end of the directory is  
36896       encountered, an application that needs to tell the difference must set *errno* to zero before the call  
36897       and check it if NULL is returned. Since the function must not change *errno* in the second case  
36898       and must set it to a non-zero value in the first case, a zero *errno* after a call returning NULL  
36899       indicates end of directory; otherwise, an error.

36900       Routines to deal with this problem more directly were proposed:

```
36901       int derror (dirp)
```

```
36902       DIR *dirp;
```

```
36903       void clearerr (dirp)
```

```
36904       DIR *dirp;
```

36905       The first would indicate whether an error had occurred, and the second would clear the error  
36906       indication. The simpler method involving *errno* was adopted instead by requiring that *readdir()*  
36907       not change *errno* when end-of-directory is encountered.

36908       An error or signal indicating that a directory has changed while open was considered but  
36909       rejected.

36910       The thread-safe version of the directory reading function returns values in a user-supplied buffer  
36911       instead of possibly using a static data area that may be overwritten by each call. Either the  
36912       {NAME\_MAX} compile-time constant or the corresponding *pathconf()* option can be used to  
36913       determine the maximum sizes of returned pathnames.

36914 **FUTURE DIRECTIONS**

36915       None.

36916 **SEE ALSO**

36917       *closedir()*, *lstat()*, *opendir()*, *rewinddir()*, *symlink()*, the Base Definitions volume of  
36918       IEEE Std 1003.1-200x, <*dirent.h*>, <*sys/types.h*>

36919 **CHANGE HISTORY**

36920       First released in Issue 2.

36921 **Issue 5**

36922       Large File Summit extensions are added.

36923       The *readdir\_r()* function is included for alignment with the POSIX Threads Extension.

36924       A note indicating that the *readdir()* function need not be reentrant is added to the  
36925       DESCRIPTION.

36926 **Issue 6**

- 36927 The `readdir_r()` function is marked as part of the Thread-Safe Functions option.
- 36928 The Open Group Corrigendum U026/7 is applied, correcting the prototype for `readdir_r()`.
- 36929 The Open Group Corrigendum U026/8 is applied, clarifying the wording of the successful  
36930 return for the `readdir_r()` function.
- 36931 The following new requirements on POSIX implementations derive from alignment with the  
36932 Single UNIX Specification:
- 36933 • The requirement to include `<sys/types.h>` has been removed. Although `<sys/types.h>` was  
36934 required for conforming implementations of previous POSIX specifications, it was not  
36935 required for UNIX applications.
  - 36936 • A statement is added to the DESCRIPTION indicating the disposition of certain fields in  
36937 **struct dirent** when an entry refers to a symbolic link.
  - 36938 • The [EOVERFLOW] mandatory error condition is added. This change is to support large  
36939 files.
  - 36940 • The [ENOENT] optional error condition is added.
- 36941 The APPLICATION USAGE section is updated to include a note on the thread-safe function and  
36942 its avoidance of possibly using a static data area.
- 36943 The **restrict** keyword is added to the `readdir_r()` prototype for alignment with the  
36944 ISO/IEC 9899:1999 standard.

36945 **NAME**

36946 readlink — read the contents of a symbolic link

36947 **SYNOPSIS**

36948 #include &lt;unistd.h&gt;

36949 ssize\_t readlink(const char \*restrict path, char \*restrict buf,  
36950 size\_t bufsize);36951 **DESCRIPTION**36952 The *readlink()* function shall place the contents of the symbolic link referred to by *path* in the  
36953 buffer *buf* which has size *bufsize*. If the number of bytes in the symbolic link is less than *bufsize*,  
36954 the contents of the remainder of *buf* are unspecified. If the *buf* argument is not large enough to  
36955 contain the link content, the first *bufsize* bytes shall be placed in *buf*.36956 If the value of *bufsize* is greater than {SSIZE\_MAX}, the result is implementation-defined.36957 **RETURN VALUE**36958 Upon successful completion, *readlink()* shall return the count of bytes placed in the buffer.  
36959 Otherwise, it shall return a value of -1, leave the buffer unchanged, and set *errno* to indicate the  
36960 error.36961 **ERRORS**36962 The *readlink()* function shall fail if:36963 [EACCES] Search permission is denied for a component of the path prefix of *path*.36964 [EINVAL] The *path* argument names a file that is not a symbolic link.

36965 [EIO] An I/O error occurred while reading from the file system.

36966 [ELOOP] A loop exists in symbolic links encountered during resolution of the *path*  
36967 argument.

36968 [ENAMETOOLONG]

36969 The length of the *path* argument exceeds {PATH\_MAX} or a pathname |  
36970 component is longer than {NAME\_MAX}. |36971 [ENOENT] A component of *path* does not name an existing file or *path* is an empty string.

36972 [ENOTDIR] A component of the path prefix is not a directory.

36973 The *readlink()* function may fail if:

36974 [EACCES] Read permission is denied for the directory.

36975 [ELOOP] More than {SYMLOOP\_MAX} symbolic links were encountered during  
36976 resolution of the *path* argument.

36977 [ENAMETOOLONG]

36978 As a result of encountering a symbolic link in resolution of the *path* argument, |  
36979 the length of the substituted pathname string exceeded {PATH\_MAX}. |

36980 **EXAMPLES**36981 **Reading the Name of a Symbolic Link**

36982 The following example shows how to read the name of a symbolic link named `/modules/pass1`.

```
36983 #include <unistd.h>
36984 char buf[1024];
36985 int len;
36986 ...
36987 if ((len = readlink("/modules/pass1", buf, sizeof(buf)-1)) != -1);
36988 buf[len] = '\0';
```

36989 **APPLICATION USAGE**

36990 Conforming applications should not assume that the returned contents of the symbolic link are null-terminated.

36992 **RATIONALE**

36993 Since IEEE Std 1003.1-200x does not require any association of file times with symbolic links, there is no requirement that file times be updated by `readlink()`. The type associated with `bufsiz` is a `size_t` in order to be consistent with both the ISO C standard and the definition of `read()`. The behavior specified for `readlink()` when `bufsiz` is zero represents historical practice. For this case, the standard developers considered a change whereby `readlink()` would return the number of non-null bytes contained in the symbolic link with the buffer `buf` remaining unchanged; however, since the `stat` structure member `st_size` value can be used to determine the size of buffer necessary to contain the contents of the symbolic link as returned by `readlink()`, this proposal was rejected, and the historical practice retained.

37002 **FUTURE DIRECTIONS**

37003 None.

37004 **SEE ALSO**

37005 `lstat()`, `stat()`, `symlink()`, the Base Definitions volume of IEEE Std 1003.1-200x, `<unistd.h>`

37006 **CHANGE HISTORY**

37007 First released in Issue 4, Version 2.

37008 **Issue 5**

37009 Moved from X/OPEN UNIX extension to BASE.

37010 **Issue 6**

37011 The return type is changed to `ssize_t`, to align with the IEEE P1003.1a draft standard.

37012 The following new requirements on POSIX implementations derive from alignment with the Single UNIX Specification:

- 37014 • This function is made mandatory.
- 37015 • In this function it is possible for the return value to exceed the range of the type `ssize_t` (since `size_t` has a larger range of positive values than `ssize_t`). A sentence restricting the size of the `size_t` object is added to the description to resolve this conflict.

37018 The following changes are made for alignment with the ISO POSIX-1:1996 standard:

- 37019 • The FUTURE DIRECTIONS section is changed to None.

37020 The following changes were made to align with the IEEE P1003.1a draft standard:

- 37021 • The [ELOOP] optional error condition is added.

37022  
37023

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

37024 **NAME**

37025 readv — read a vector

37026 **SYNOPSIS**

37027 XSI #include &lt;sys/uio.h&gt;

37028 ssize\_t readv(int *fildes*, const struct iovec \**iovcnt*, int *iovcnt*);

37029

37030 **DESCRIPTION**

37031 The *readv()* function shall be equivalent to *read()*, except as described below. The *readv()*  
 37032 function shall place the input data into the *iovcnt* buffers specified by the members of the *iovcnt*  
 37033 array: *iovcnt*[0], *iovcnt*[1], ..., *iovcnt*[*iovcnt*-1]. The *iovcnt* argument is valid if greater than 0 and less than  
 37034 or equal to {IOV\_MAX}.

37035 Each *iovcnt* entry specifies the base address and length of an area in memory where data should  
 37036 be placed. The *readv()* function shall always fill an area completely before proceeding to the  
 37037 next.

37038 Upon successful completion, *readv()* shall mark for update the *st\_atime* field of the file.

37039 **RETURN VALUE**37040 Refer to *read()*.37041 **ERRORS**37042 Refer to *read()*.37043 In addition, the *readv()* function shall fail if:37044 [EINVAL] The sum of the *iovcnt* values in the *iovcnt* array overflowed an *ssize\_t*.37045 The *readv()* function may fail if:37046 [EINVAL] The *iovcnt* argument was less than or equal to 0, or greater than {IOV\_MAX}.37047 **EXAMPLES**37048 **Reading Data into an Array**

37049 The following example reads data from the file associated with the file descriptor *fd* into the  
 37050 buffers specified by members of the *iovcnt* array.

37051 #include &lt;sys/types.h&gt;

37052 #include &lt;sys/uio.h&gt;

37053 #include &lt;unistd.h&gt;

37054 ...

37055 ssize\_t bytes\_read;

37056 int fd;

37057 char buf0[20];

37058 char buf1[30];

37059 char buf2[40];

37060 int iovcnt;

37061 struct iovec iovcnt[3];

37062 iovcnt[0].iovcnt\_base = buf0;

37063 iovcnt[0].iovcnt\_len = sizeof(buf0);

37064 iovcnt[1].iovcnt\_base = buf1;

37065 iovcnt[1].iovcnt\_len = sizeof(buf1);

37066 iovcnt[2].iovcnt\_base = buf2;

37067 iovcnt[2].iovcnt\_len = sizeof(buf2);

```
37068     ...
37069     iovcnt = sizeof(iov) / sizeof(struct iovec);
37070     bytes_read = readv(fd, iov, iovcnt);
37071     ...

37072 APPLICATION USAGE
37073     None.

37074 RATIONALE
37075     Refer to read().

37076 FUTURE DIRECTIONS
37077     None.

37078 SEE ALSO
37079     read(), writew(), the Base Definitions volume of IEEE Std 1003.1-200x, <sys/uio.h>

37080 CHANGE HISTORY
37081     First released in Issue 4, Version 2.

37082 Issue 6
37083     Split out from the read() reference page.
```

37084 **NAME**

37085        realloc — memory reallocator

37086 **SYNOPSIS**

37087        #include &lt;stdlib.h&gt;

37088        void \*realloc(void \*ptr, size\_t size);

37089 **DESCRIPTION**

37090 **CX**        The functionality described on this reference page is aligned with the ISO C standard. Any  
37091 conflict between the requirements described here and the ISO C standard is unintentional. This  
37092 volume of IEEE Std 1003.1-200x defers to the ISO C standard.

37093        The *realloc()* function shall change the size of the memory object pointed to by *ptr* to the size  
37094 specified by *size*. The contents of the object shall remain unchanged up to the lesser of the new  
37095 and old sizes. If the new size of the memory object would require movement of the object, the  
37096 space for the previous instantiation of the object is freed. If the new size is larger, the contents of  
37097 the newly allocated portion of the object are unspecified. If *size* is 0 and *ptr* is not a null pointer,  
37098 the object pointed to is freed. If the space cannot be allocated, the object shall remain unchanged.

37099        If *ptr* is a null pointer, *realloc()* shall be equivalent to *malloc()* for the specified size.

37100        If *ptr* does not match a pointer returned earlier by *calloc()*, *malloc()*, or *realloc()* or if the space has  
37101 previously been deallocated by a call to *free()* or *realloc()*, the behavior is undefined.

37102        The order and contiguity of storage allocated by successive calls to *realloc()* is unspecified. The  
37103 pointer returned if the allocation succeeds shall be suitably aligned so that it may be assigned to  
37104 a pointer to any type of object and then used to access such an object in the space allocated (until  
37105 the space is explicitly freed or reallocated). Each such allocation shall yield a pointer to an object  
37106 disjoint from any other object. The pointer returned shall point to the start (lowest byte address)  
37107 of the allocated space. If the space cannot be allocated, a null pointer shall be returned.

37108 **RETURN VALUE**

37109        Upon successful completion with a *size* not equal to 0, *realloc()* shall return a pointer to the  
37110 (possibly moved) allocated space. If *size* is 0, either a null pointer or a unique pointer that can be  
37111 successfully passed to *free()* shall be returned. If there is not enough available memory, *realloc()*  
37112 **CX** shall return a null pointer and set *errno* to [ENOMEM].

37113 **ERRORS**37114        The *realloc()* function shall fail if:

37115 **CX**        [ENOMEM]        Insufficient memory is available.

37116 **EXAMPLES**

37117        None.

37118 **APPLICATION USAGE**

37119        None.

37120 **RATIONALE**

37121        None.

37122 **FUTURE DIRECTIONS**

37123        None.

37124 **SEE ALSO**37125        *calloc()*, *free()*, *malloc()*, the Base Definitions volume of IEEE Std 1003.1-200x, <stdlib.h>



37126 **CHANGE HISTORY**

37127 First released in Issue 1. Derived from Issue 1 of the SVID.

37128 **Issue 6**

37129 Extensions beyond the ISO C standard are now marked.

37130 The following new requirements on POSIX implementations derive from alignment with the  
37131 Single UNIX Specification:

- 37132 • In the RETURN VALUE section, if there is not enough available memory, the setting of *errno*  
37133 to [ENOMEM] is added.
- 37134 • The [ENOMEM] error condition is added.

## 37135 NAME

37136        realpath — resolve a pathname |

## 37137 SYNOPSIS

37138 xSI        #include &lt;stdlib.h&gt;

37139        char \*realpath(const char \*restrict *file\_name*,  
37140                      char \*restrict *resolved\_name*);

37141

## 37142 DESCRIPTION

37143        The *realpath()* function shall derive, from the pathname pointed to by *file\_name*, an absolute |  
 37144        pathname that names the same file, whose resolution does not involve '.', '..', or symbolic |  
 37145        links. The generated pathname shall be stored as a null-terminated string, up to a maximum of |  
 37146        {PATH\_MAX} bytes, in the buffer pointed to by *resolved\_name*.

## 37147 RETURN VALUE

37148        Upon successful completion, *realpath()* shall return a pointer to the resolved name. Otherwise,  
 37149        *realpath()* shall return a null pointer and set *errno* to indicate the error, and the contents of the  
 37150        buffer pointed to by *resolved\_name* are undefined.

## 37151 ERRORS

37152        The *realpath()* function shall fail if:37153        [EACCES]        Read or search permission was denied for a component of *file\_name*.37154        [EINVAL]        Either the *file\_name* or *resolved\_name* argument is a null pointer.

37155        [EIO]            An error occurred while reading from the file system.

37156        [ELOOP]        A loop exists in symbolic links encountered during resolution of the *path*  
37157        argument.

37158        [ENAMETOOLONG]

37159                      The length of the *file\_name* argument exceeds {PATH\_MAX} or a pathname |  
37160                      component is longer than {NAME\_MAX}. |37161        [ENOENT]        A component of *file\_name* does not name an existing file or *file\_name* points to  
37162        an empty string.

37163        [ENOTDIR]       A component of the path prefix is not a directory.

37164        The *realpath()* function may fail if:37165        [ELOOP]        More than {SYMLOOP\_MAX} symbolic links were encountered during  
37166        resolution of the *path* argument.

37167        [ENAMETOOLONG]

37168                      Pathname resolution of a symbolic link produced an intermediate result |  
37169                      whose length exceeds {PATH\_MAX}. |

37170        [ENOMEM]        Insufficient storage space is available.

37171 **EXAMPLES**37172 **Generating an Absolute Pathname** |

37173 The following example generates an absolute pathname for the file identified by the *symlinkpath* |  
37174 argument. The generated pathname is stored in the *actualpath* array. |

```
37175 #include <stdlib.h>
37176 ...
37177 char *symlinkpath = "/tmp/symlink/file";
37178 char actualpath [PATH_MAX+1];
37179 char *ptr;

37180 ptr = realpath(symlinkpath, actualpath);
```

37181 **APPLICATION USAGE**

37182 None.

37183 **RATIONALE**

37184 None.

37185 **FUTURE DIRECTIONS**

37186 None.

37187 **SEE ALSO**

37188 *getcwd()*, *sysconf()*, the Base Definitions volume of IEEE Std 1003.1-200x, <**stdlib.h**>

37189 **CHANGE HISTORY**

37190 First released in Issue 4, Version 2.

37191 **Issue 5**

37192 Moved from X/OPEN UNIX extension to BASE.

37193 **Issue 6**

37194 The **restrict** keyword is added to the *realpath()* prototype for alignment with the |  
37195 ISO/IEC 9899:1999 standard.

37196 The wording of the mandatory [ELOOP] error condition is updated, and a second optional  
37197 [ELOOP] error condition is added.

37198 **NAME**

37199       recv — receive a message from a connected socket

37200 **SYNOPSIS**

37201       #include &lt;sys/socket.h&gt;

37202       ssize\_t recv(int *socket*, void \**buffer*, size\_t *length*, int *flags*);37203 **DESCRIPTION**

37204       The *recv()* function shall receive a message from a connection-mode or connectionless-mode  
 37205       socket. It is normally used with connected sockets because it does not permit the application to  
 37206       retrieve the source address of received data.

37207       The *recv()* function takes the following arguments:37208       *socket*       Specifies the socket file descriptor.37209       *buffer*       Points to a buffer where the message should be stored.37210       *length*       Specifies the length in bytes of the buffer pointed to by the *buffer* argument.

37211       *flags*       Specifies the type of message reception. Values of this argument are formed by  
 37212       logically OR'ing zero or more of the following values:

37213               MSG\_PEEK       Peeks at an incoming message. The data is treated as unread and  
 37214               the next *recv()* or similar function shall still return this data.

37215               MSG\_OOB       Requests out-of-band data. The significance and semantics of  
 37216               out-of-band data are protocol-specific.

37217               MSG\_WAITALL   On SOCK\_STREAM sockets this requests that the function block |  
 37218               until the full amount of data can be returned. The function may |  
 37219               return the smaller amount of data if the socket is a message- |  
 37220               based socket, if a signal is caught, if the connection is |  
 37221               terminated, if MSG\_PEEK was specified, or if an error is pending |  
 37222               for the socket. |

37223       The *recv()* function shall return the length of the message written to the buffer pointed to by the  
 37224       *buffer* argument. For message-based sockets, such as SOCK\_DGRAM and SOCK\_SEQPACKET,  
 37225       the entire message shall be read in a single operation. If a message is too long to fit in the  
 37226       supplied buffer, and MSG\_PEEK is not set in the *flags* argument, the excess bytes shall be  
 37227       discarded. For stream-based sockets, such as SOCK\_STREAM, message boundaries shall be  
 37228       ignored. In this case, data shall be returned to the user as soon as it becomes available, and no  
 37229       data shall be discarded. |

37230       If the MSG\_WAITALL flag is not set, data shall be returned only up to the end of the first  
 37231       message.

37232       If no messages are available at the socket and O\_NONBLOCK is not set on the socket's file  
 37233       descriptor, *recv()* shall block until a message arrives. If no messages are available at the socket  
 37234       and O\_NONBLOCK is set on the socket's file descriptor, *recv()* shall fail and set *errno* to  
 37235       [EAGAIN] or [EWOULDBLOCK].

37236 **RETURN VALUE**

37237       Upon successful completion, *recv()* shall return the length of the message in bytes. If no  
 37238       messages are available to be received and the peer has performed an orderly shutdown, *recv()*  
 37239       shall return 0. Otherwise, -1 shall be returned and *errno* set to indicate the error.

37240 **ERRORS**

- 37241 The *recv()* function shall fail if:
- 37242 [EAGAIN] or [EWOULDBLOCK]
- 37243 The socket's file descriptor is marked O\_NONBLOCK and no data is waiting  
37244 to be received; or MSG\_OOB is set and no out-of-band data is available and  
37245 either the socket's file descriptor is marked O\_NONBLOCK or the socket does  
37246 not support blocking to await out-of-band data.
- 37247 [EBADF] The *socket* argument is not a valid file descriptor.
- 37248 [ECONNRESET] A connection was forcibly closed by a peer.
- 37249 [EINTR] The *recv()* function was interrupted by a signal that was caught, before any  
37250 data was available.
- 37251 [EINVAL] The MSG\_OOB flag is set and no out-of-band data is available.
- 37252 [ENOTCONN] A receive is attempted on a connection-mode socket that is not connected.
- 37253 [ENOTSOCK] The *socket* argument does not refer to a socket.
- 37254 [EOPNOTSUPP] The specified flags are not supported for this socket type or protocol.
- 37255 [ETIMEDOUT] The connection timed out during connection establishment, or due to a  
37256 transmission timeout on active connection.
- 37257 The *recv()* function may fail if:
- 37258 [EIO] An I/O error occurred while reading from or writing to the file system.
- 37259 [ENOBUFS] Insufficient resources were available in the system to perform the operation.
- 37260 [ENOMEM] Insufficient memory was available to fulfill the request.

37261 **EXAMPLES**

37262 None.

37263 **APPLICATION USAGE**

37264 The *recv()* function is equivalent to *recvfrom()* with a zero *address\_len* argument, and to *read()* if  
37265 no flags are used.

37266 The *select()* and *poll()* functions can be used to determine when data is available to be received.

37267 **RATIONALE**

37268 None.

37269 **FUTURE DIRECTIONS**

37270 None.

37271 **SEE ALSO**

37272 *poll()*, *read()*, *recvmsg()*, *recvfrom()*, *select()*, *send()*, *sendmsg()*, *sendto()*, *shutdown()*, *socket()*,  
37273 *write()*, the Base Definitions volume of IEEE Std 1003.1-200x, <sys/socket.h>

37274 **CHANGE HISTORY**

37275 First released in Issue 6. Derived from the XNS, Issue 5.2 specification.

## 37276 NAME

37277 recvfrom — receive a message from a socket

## 37278 SYNOPSIS

37279 #include &lt;sys/socket.h&gt;

37280 ssize\_t recvfrom(int socket, void \*restrict buffer, size\_t length,

37281 int flags, struct sockaddr \*restrict address,

37282 socklen\_t \*restrict address\_len);

## 37283 DESCRIPTION

37284 The *recvfrom()* function shall receive a message from a connection-mode or connectionless-mode  
 37285 socket. It is normally used with connectionless-mode sockets because it permits the application  
 37286 to retrieve the source address of received data.

37287 The *recvfrom()* function takes the following arguments:

37288	<i>socket</i>	Specifies the socket file descriptor.
37289	<i>buffer</i>	Points to the buffer where the message should be stored.
37290	<i>length</i>	Specifies the length in bytes of the buffer pointed to by the <i>buffer</i> argument.
37291	<i>flags</i>	Specifies the type of message reception. Values of this argument are formed 37292 by logically OR'ing zero or more of the following values:
37293	MSG_PEEK	Peeks at an incoming message. The data is treated as unread 37294 and the next <i>recvfrom()</i> or similar function shall still return 37295 this data.
37296	MSG_OOB	Requests out-of-band data. The significance and semantics 37297 of out-of-band data are protocol-specific.
37298	MSG_WAITALL	On SOCK_STREAM sockets this requests that the function   37299 block until the full amount of data can be returned. The   37300 function may return the smaller amount of data if the socket   37301 is a message-based socket, if a signal is caught, if the   37302 connection is terminated, if MSG_PEEK was specified, or if   37303 an error is pending for the socket.
37304	<i>address</i>	A null pointer, or points to a <b>sockaddr</b> structure in which the sending address 37305 is to be stored. The length and format of the address depend on the address 37306 family of the socket.
37307	<i>address_len</i>	Specifies the length of the <b>sockaddr</b> structure pointed to by the <i>address</i> 37308 argument.

37309 The *recvfrom()* function shall return the length of the message written to the buffer pointed to by  
 37310 RS the *buffer* argument. For message-based sockets, such as SOCK\_RAW, SOCK\_DGRAM, and  
 37311 SOCK\_SEQPACKET, the entire message shall be read in a single operation. If a message is too  
 37312 long to fit in the supplied buffer, and MSG\_PEEK is not set in the *flags* argument, the excess  
 37313 bytes shall be discarded. For stream-based sockets, such as SOCK\_STREAM, message  
 37314 boundaries shall be ignored. In this case, data shall be returned to the user as soon as it becomes |  
 37315 available, and no data shall be discarded.

37316 If the MSG\_WAITALL flag is not set, data shall be returned only up to the end of the first  
 37317 message.

37318 Not all protocols provide the source address for messages. If the *address* argument is not a null  
 37319 pointer and the protocol provides the source address of messages, the source address of the |

37320 received message shall be stored in the **sockaddr** structure pointed to by the *address* argument, |  
 37321 and the length of this address shall be stored in the object pointed to by the *address\_len* |  
 37322 argument.

37323 If the actual length of the address is greater than the length of the supplied **sockaddr** structure, |  
 37324 the stored address shall be truncated.

37325 If the *address* argument is not a null pointer and the protocol does not provide the source address |  
 37326 of messages, the value stored in the object pointed to by *address* is unspecified.

37327 If no messages are available at the socket and O\_NONBLOCK is not set on the socket's file |  
 37328 descriptor, *recvfrom()* shall block until a message arrives. If no messages are available at the |  
 37329 socket and O\_NONBLOCK is set on the socket's file descriptor, *recvfrom()* shall fail and set *errno* |  
 37330 to [EAGAIN] or [EWOULDBLOCK].

### 37331 RETURN VALUE

37332 Upon successful completion, *recvfrom()* shall return the length of the message in bytes. If no |  
 37333 messages are available to be received and the peer has performed an orderly shutdown, |  
 37334 *recvfrom()* shall return 0. Otherwise, the function shall return -1 and set *errno* to indicate the |  
 37335 error.

### 37336 ERRORS

37337 The *recvfrom()* function shall fail if:

37338 [EAGAIN] or [EWOULDBLOCK]

37339 The socket's file descriptor is marked O\_NONBLOCK and no data is waiting |  
 37340 to be received; or MSG\_OOB is set and no out-of-band data is available and |  
 37341 either the socket's file descriptor is marked O\_NONBLOCK or the socket does |  
 37342 not support blocking to await out-of-band data.

37343 [EBADF] The *socket* argument is not a valid file descriptor.

37344 [ECONNRESET] A connection was forcibly closed by a peer.

37345 [EINTR] A signal interrupted *recvfrom()* before any data was available.

37346 [EINVAL] The MSG\_OOB flag is set and no out-of-band data is available.

37347 [ENOTCONN] A receive is attempted on a connection-mode socket that is not connected.

37348 [ENOTSOCK] The *socket* argument does not refer to a socket.

37349 [EOPNOTSUPP] The specified flags are not supported for this socket type.

37350 [ETIMEDOUT] The connection timed out during connection establishment, or due to a |  
 37351 transmission timeout on active connection.

37352 The *recvfrom()* function may fail if:

37353 [EIO] An I/O error occurred while reading from or writing to the file system.

37354 [ENOBUFS] Insufficient resources were available in the system to perform the operation.

37355 [ENOMEM] Insufficient memory was available to fulfill the request.

37356 **EXAMPLES**

37357 None.

37358 **APPLICATION USAGE**37359 The *select()* and *poll()* functions can be used to determine when data is available to be received.37360 **RATIONALE**

37361 None.

37362 **FUTURE DIRECTIONS**

37363 None.

37364 **SEE ALSO**37365 *poll()*, *read()*, *recv()*, *recvmsg()*, *select()* (on page 1742)1 *send()*, *sendmsg()*, *sendto()*, *shutdown()*,37366 *socket()*, *write()*, the Base Definitions volume of IEEE Std 1003.1-200x, <sys/socket.h>37367 **CHANGE HISTORY**

37368 First released in Issue 6. Derived from the XNS, Issue 5.2 specification.



37369 **NAME**

37370 recvmsg — receive a message from a socket

37371 **SYNOPSIS**

37372 #include &lt;sys/socket.h&gt;

37373 ssize\_t recvmsg(int *socket*, struct msghdr \**message*, int *flags*);37374 **DESCRIPTION**

37375 The *recvmsg()* function shall receive a message from a connection-mode or connectionless-mode  
 37376 socket. It is normally used with connectionless-mode sockets because it permits the application  
 37377 to retrieve the source address of received data.

37378 The *recvmsg()* function takes the following arguments:

37379	<i>socket</i>	Specifies the socket file descriptor.
37380	<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.
37381		
37382		
37383		
37384	<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:
37385		
37386	MSG_OOB	Requests out-of-band data. The significance and semantics of out-of-band data are protocol-specific.
37387		
37388	MSG_PEEK	Peeks at the incoming message.
37389	MSG_WAITALL	On SOCK_STREAM sockets this requests that the function
37390		block until the full amount of data can be returned. The
37391		function may return the smaller amount of data if the socket
37392		is a message-based socket, if a signal is caught, if the
37393		connection is terminated, if MSG_PEEK was specified, or if
37394		an error is pending for the socket.

37395 The *recvmsg()* function shall receive messages from unconnected or connected sockets and shall  
 37396 return the length of the message.

37397 The *recvmsg()* function shall return the total length of the message. For message-based sockets,  
 37398 such as SOCK\_DGRAM and SOCK\_SEQPACKET, the entire message shall be read in a single  
 37399 operation. If a message is too long to fit in the supplied buffers, and MSG\_PEEK is not set in the  
 37400 *flags* argument, the excess bytes shall be discarded, and MSG\_TRUNC shall be set in the  
 37401 *msg\_flags* member of the **msghdr** structure. For stream-based sockets, such as SOCK\_STREAM,  
 37402 message boundaries shall be ignored. In this case, data shall be returned to the user as soon as it  
 37403 becomes available, and no data shall be discarded.

37404 If the MSG\_WAITALL flag is not set, data shall be returned only up to the end of the first  
 37405 message.

37406 If no messages are available at the socket and O\_NONBLOCK is not set on the socket's file  
 37407 descriptor, *recvmsg()* shall block until a message arrives. If no messages are available at the  
 37408 socket and O\_NONBLOCK is set on the socket's file descriptor, *recvmsg()* function shall fail and  
 37409 set *errno* to [EAGAIN] or [EWOULDBLOCK].

37410 In the **msghdr** structure, the *msg\_name* and *msg\_namelen* members specify the source address if  
 37411 the socket is unconnected. If the socket is connected, the *msg\_name* and *msg\_namelen* members  
 37412 shall be ignored. The *msg\_name* member may be a null pointer if no names are desired or  
 37413 required. The *msg\_iov* and *msg\_iovlen* fields are used to specify where the received data shall be

37414 stored. *msg\_iov* points to an array of **iovec** structures; *msg\_iovlen* shall be set to the dimension of  
 37415 this array. In each **iovec** structure, the *iov\_base* field specifies a storage area and the *iov\_len* field  
 37416 gives its size in bytes. Each storage area indicated by *msg\_iov* is filled with received data in turn  
 37417 until all of the received data is stored or all of the areas have been filled.

37418 Upon successful completion, the *msg\_flags* member of the message header shall be the bitwise- |  
 37419 inclusive OR of all of the following flags that indicate conditions detected for the received |  
 37420 message: |

37421 MSG\_EOR End of record was received (if supported by the protocol).

37422 MSG\_OOB Out-of-band data was received.

37423 MSG\_TRUNC Normal data was truncated.

37424 MSG\_CTRUNC Control data was truncated.

#### 37425 RETURN VALUE

37426 Upon successful completion, *recvmsg()* shall return the length of the message in bytes. If no  
 37427 messages are available to be received and the peer has performed an orderly shutdown,  
 37428 *recvmsg()* shall return 0. Otherwise,  $-1$  shall be returned and *errno* set to indicate the error.

#### 37429 ERRORS

37430 The *recvmsg()* function shall fail if:

37431 [EAGAIN] or [EWOULDBLOCK]

37432 The socket's file descriptor is marked O\_NONBLOCK and no data is waiting  
 37433 to be received; or MSG\_OOB is set and no out-of-band data is available and  
 37434 either the socket's file descriptor is marked O\_NONBLOCK or the socket does  
 37435 not support blocking to await out-of-band data.

37436 [EBADF] The *socket* argument is not a valid open file descriptor.

37437 [ECONNRESET] A connection was forcibly closed by a peer.

37438 [EINTR] This function was interrupted by a signal before any data was available.

37439 [EINVAL] The sum of the *iov\_len* values overflows a **ssize\_t**, or the MSG\_OOB flag is set  
 37440 and no out-of-band data is available.

37441 [EMSGSIZE] The *msg\_iovlen* member of the **msghdr** structure pointed to by *message* is less  
 37442 than or equal to 0, or is greater than {IOV\_MAX}.

37443 [ENOTCONN] A receive is attempted on a connection-mode socket that is not connected.

37444 [ENOTSOCK] The *socket* argument does not refer to a socket.

37445 [EOPNOTSUPP] The specified flags are not supported for this socket type.

37446 [ETIMEDOUT] The connection timed out during connection establishment, or due to a  
 37447 transmission timeout on active connection.

37448 The *recvmsg()* function may fail if:

37449 [EIO] An I/O error occurred while reading from or writing to the file system.

37450 [ENOBUFS] Insufficient resources were available in the system to perform the operation.

37451 [ENOMEM] Insufficient memory was available to fulfill the request.

37452 **EXAMPLES**

37453 None.

37454 **APPLICATION USAGE**37455 The *select()* and *poll()* functions can be used to determine when data is available to be received.37456 **RATIONALE**

37457 None.

37458 **FUTURE DIRECTIONS**

37459 None.

37460 **SEE ALSO**37461 *poll()*, *recv()*, *recvfrom()*, *select()*, *send()*, *sendmsg()*, *sendto()*, *shutdown()*, *socket()*, the Base

37462 Definitions volume of IEEE Std 1003.1-200x, &lt;sys/socket.h&gt;

37463 **CHANGE HISTORY**

37464 First released in Issue 6. Derived from the XNS, Issue 5.2 specification.

37465 **NAME**

37466 regcomp, regerror, regex, regfree — regular expression matching

37467 **SYNOPSIS**

37468 #include <regex.h>

```

37469 int regcomp(regex_t *restrict preg, const char *restrict pattern, int cflags);
37470 size_t regerror(int errcode, const regex_t *restrict preg,
37471 char *restrict errbuf, size_t errbuf_size);
37472 int regexexec(const regex_t *restrict preg, const char *restrict string,
37473 size_t rmatch, regmatch_t pmatch[restrict], int eflags);
37474 void regfree(regex_t *preg);
    
```

37475 **DESCRIPTION**

37476 These functions interpret *basic* and *extended* regular expressions as described in the Base  
 37477 Definitions volume of IEEE Std 1003.1-200x, Chapter 9, Regular Expressions.

37478 The **regex\_t** structure is defined in <regex.h> and contains at least the following member: |

37479  
 37480  
 37481

Member Type	Member Name	Description
size_t	re_nsub	Number of parenthesized subexpressions.

37482 The **regmatch\_t** structure is defined in <regex.h> and contains at least the following members: |

37483  
 37484  
 37485  
 37486  
 37487

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.

37488 The *regcomp()* function shall compile the regular expression contained in the string pointed to by  
 37489 the *pattern* argument and place the results in the structure pointed to by *preg*. The *cflags*  
 37490 argument is the bitwise-inclusive OR of zero or more of the following flags, which are defined in  
 37491 the <regex.h> header: |

- 37492 REG\_EXTENDED     Use Extended Regular Expressions.
- 37493 REG\_ICASE        Ignore case in match. (See the Base Definitions volume of  
 37494 IEEE Std 1003.1-200x, Chapter 9, Regular Expressions.)
- 37495 REG\_NOSUB        Report only success/fail in *regexexec()*.
- 37496 REG\_NEWLINE     Change the handling of <newline>s, as described in the text.

37497 The default regular expression type for *pattern* is a Basic Regular Expression. The application can  
 37498 specify Extended Regular Expressions using the REG\_EXTENDED *cflags* flag.

37499 If the REG\_NOSUB flag was not set in *cflags*, then *regcomp()* shall set *re\_nsub* to the number of  
 37500 parenthesized subexpressions (delimited by "\(\)" in basic regular expressions or "( )" in  
 37501 extended regular expressions) found in *pattern*.

37502 The *regexexec()* function compares the null-terminated string specified by *string* with the compiled  
 37503 regular expression *preg* initialized by a previous call to *regcomp()*. If it finds a match, *regexexec()*  
 37504 shall return 0; otherwise, it shall return non-zero indicating either no match or an error. The  
 37505 *eflags* argument is the bitwise-inclusive OR of zero or more of the following flags, which are  
 37506 defined in the <regex.h> header:

37507 REG\_NOTBOL The first character of the string pointed to by *string* is not the beginning of the  
 37508 line. Therefore, the circumflex character ('^'), when taken as a special  
 37509 character, shall not match the beginning of *string*.

37510 REG\_NOTEOL The last character of the string pointed to by *string* is not the end of the line.  
 37511 Therefore, the dollar sign ('\$'), when taken as a special character, shall not  
 37512 match the end of *string*.

37513 If *nmatch* is 0 or REG\_NOSUB was set in the *cflags* argument to *regcomp()*, then *regexec()* shall  
 37514 ignore the *pmatch* argument. Otherwise, the application shall ensure that the *pmatch* argument  
 37515 points to an array with at least *nmatch* elements, and *regexec()* shall fill in the elements of that  
 37516 array with offsets of the substrings of *string* that correspond to the parenthesized subexpressions  
 37517 of *pattern*: *pmatch[i].rm\_so* shall be the byte offset of the beginning and *pmatch[i].rm\_eo* shall be  
 37518 one greater than the byte offset of the end of substring *i*. (Subexpression *i* begins at the *i*th  
 37519 matched open parenthesis, counting from 1.) Offsets in *pmatch[0]* identify the substring that  
 37520 corresponds to the entire regular expression. Unused elements of *pmatch* up to *pmatch[nmatch-1]*  
 37521 shall be filled with -1. If there are more than *nmatch* subexpressions in *pattern* (*pattern* itself  
 37522 counts as a subexpression), then *regexec()* shall still do the match, but shall record only the first  
 37523 *nmatch* substrings.

37524 When matching a basic or extended regular expression, any given parenthesized subexpression  
 37525 of *pattern* might participate in the match of several different substrings of *string*, or it might not  
 37526 match any substring even though the pattern as a whole did match. The following rules shall be  
 37527 used to determine which substrings to report in *pmatch* when matching regular expressions:

37528 1. If subexpression *i* in a regular expression is not contained within another subexpression,  
 37529 and it participated in the match several times, then the byte offsets in *pmatch[i]* shall  
 37530 delimit the last such match.

37531 2. If subexpression *i* is not contained within another subexpression, and it did not participate  
 37532 in an otherwise successful match, the byte offsets in *pmatch[i]* shall be -1. A subexpression  
 37533 does not participate in the match when:

37534 ' \* ' or "\{\}" appears immediately after the subexpression in a basic regular  
 37535 expression, or ' \* ', '? ', or "{ }" appears immediately after the subexpression in an  
 37536 extended regular expression, and the subexpression did not match (matched 0 times)

37537 or:

37538 ' | ' is used in an extended regular expression to select this subexpression or another,  
 37539 and the other subexpression matched.

37540 3. If subexpression *i* is contained within another subexpression *j*, and *i* is not contained  
 37541 within any other subexpression that is contained within *j*, and a match of subexpression *j*  
 37542 is reported in *pmatch[j]*, then the match or non-match of subexpression *i* reported in  
 37543 *pmatch[i]* shall be as described in 1. and 2. above, but within the substring reported in  
 37544 *pmatch[j]* rather than the whole string. The offsets in *pmatch[i]* are still relative to the start  
 37545 of string.

37546 4. If subexpression *i* is contained in subexpression *j*, and the byte offsets in *pmatch[j]* are -1,  
 37547 then the pointers in *pmatch[i]* shall also be -1.

37548 5. If subexpression *i* matched a zero-length string, then both byte offsets in *pmatch[i]* shall be  
 37549 the byte offset of the character or null terminator immediately following the zero-length  
 37550 string.

37551 If, when *regexec()* is called, the locale is different from when the regular expression was  
 37552 compiled, the result is undefined.

37553 If REG\_NEWLINE is not set in *cflags*, then a <newline> in *pattern* or *string* shall be treated as an  
 37554 ordinary character. If REG\_NEWLINE is set, then <newline> shall be treated as an ordinary  
 37555 character except as follows:

- 37556 1. A <newline> in *string* shall not be matched by a period outside a bracket expression or by  
 37557 any form of a non-matching list (see the Base Definitions volume of IEEE Std 1003.1-200x,  
 37558 Chapter 9, Regular Expressions).
- 37559 2. A circumflex ('^') in *pattern*, when used to specify expression anchoring (see the Base  
 37560 Definitions volume of IEEE Std 1003.1-200x, Section 9.3.8, BRE Expression Anchoring),  
 37561 shall match the zero-length string immediately after a <newline> in *string*, regardless of  
 37562 the setting of REG\_NOTBOL.
- 37563 3. A dollar sign ('\$') in *pattern*, when used to specify expression anchoring, shall match the  
 37564 zero-length string immediately before a <newline> in *string*, regardless of the setting of  
 37565 REG\_NOTEOL.

37566 The *regfree()* function frees any memory allocated by *regcomp()* associated with *preg*.

37567 The following constants are defined as error return values:

37568	REG_NOMATCH	<i>regexec()</i> failed to match.
37569	REG_BADPAT	Invalid regular expression.
37570	REG_ECOLLATE	Invalid collating element referenced.
37571	REG_ECTYPE	Invalid character class type referenced.
37572	REG_EESCAPE	Trailing '\\' in pattern.
37573	REG_ESUBREG	Number in "\digit" invalid or in error.
37574	REG_EBRACK	"[]" imbalance.
37575	REG_EPAREN	"\(\)" or "()" imbalance.
37576	REG_EBRACE	"\{\}" imbalance.
37577	REG_BADBR	Content of "\{\}" invalid: not a number, number too large, more than 37578 two numbers, first larger than second.
37579	REG_ERANGE	Invalid endpoint in range expression.
37580	REG_ESPACE	Out of memory.
37581	REG_BADRPT	'?', '*', or '+' not preceded by valid regular expression.

37582 The *regerror()* function provides a mapping from error codes returned by *regcomp()* and  
 37583 *regexec()* to unspecified printable strings. It generates a string corresponding to the value of the  
 37584 *errcode* argument, which the application shall ensure is the last non-zero value returned by  
 37585 *regcomp()* or *regexec()* with the given value of *preg*. If *errcode* is not such a value, the content of  
 37586 the generated string is unspecified.

37587 If *preg* is a null pointer, but *errcode* is a value returned by a previous call to *regexec()* or *regcomp()*,  
 37588 the *regerror()* still generates an error string corresponding to the value of *errcode*, but it might not  
 37589 be as detailed under some implementations.

37590 If the *errbuf\_size* argument is not 0, *regerror()* shall place the generated string into the buffer of  
 37591 size *errbuf\_size* bytes pointed to by *errbuf*. If the string (including the terminating null) cannot fit  
 37592 in the buffer, *regerror()* shall truncate the string and null-terminates the result.

37593 If *errbuf\_size* is 0, *regerror()* shall ignore the *errbuf* argument, and return the size of the buffer  
 37594 needed to hold the generated string.

37595 If the *preg* argument to *regexec()* or *regfree()* is not a compiled regular expression returned by  
 37596 *regcomp()*, the result is undefined. A *preg* is no longer treated as a compiled regular expression  
 37597 after it is given to *regfree()*.

#### 37598 RETURN VALUE

37599 Upon successful completion, the *regcomp()* function shall return 0. Otherwise, it shall return an  
 37600 integer value indicating an error as described in <**regex.h**>, and the content of *preg* is undefined.  
 37601 If a code is returned, the interpretation shall be as given in <**regex.h**>.

37602 If *regcomp()* detects an invalid RE, it may return REG\_BADPAT, or it may return one of the error  
 37603 codes that more precisely describes the error.

37604 Upon successful completion, the *regexec()* function shall return 0. Otherwise, it shall return  
 37605 REG\_NOMATCH to indicate no match.

37606 Upon successful completion, the *regerror()* function shall return the number of bytes needed to  
 37607 hold the entire generated string, including the null termination. If the return value is greater than  
 37608 *errbuf\_size*, the string returned in the buffer pointed to by *errbuf* has been truncated.

37609 The *regfree()* function shall not return a value.

#### 37610 ERRORS

37611 No errors are defined.

#### 37612 EXAMPLES

```

37613 #include <regex.h>
37614 /*
37615  * Match string against the extended regular expression in
37616  * pattern, treating errors as no match.
37617  *
37618  * Return 1 for match, 0 for no match.
37619  */
37620 int
37621 match(const char *string, char *pattern)
37622 {
37623     int    status;
37624     regex_t re;
37625     if (regcomp(&re, pattern, REG_EXTENDED|REG_NOSUB) != 0) {
37626         return(0); /* Report error. */
37627     }
37628     status = regexec(&re, string, (size_t) 0, NULL, 0);
37629     regfree(&re);
37630     if (status != 0) {
37631         return(0); /* Report error. */
37632     }
37633     return(1);
37634 }

```

37635 The following demonstrates how the REG\_NOTBOL flag could be used with *regexec()* to find all  
 37636 substrings in a line that match a pattern supplied by a user. (For simplicity of the example, very  
 37637 little error checking is done.)

```

37638     (void) regcomp (&re, pattern, 0);
37639     /* This call to regexec() finds the first match on the line. */
37640     error = regexec (&re, &buffer[0], 1, &pm, 0);
37641     while (error == 0) { /* While matches found. */
37642         /* Substring found between pm.rm_so and pm.rm_eo. */
37643         /* This call to regexec() finds the next match. */
37644         error = regexec (&re, buffer + pm.rm_eo, 1, &pm, REG_NOTBOL);
37645     }

```

#### 37646 APPLICATION USAGE

37647 An application could use:

```
37648     regerror(code, preg, (char *)NULL, (size_t)0)
```

37649 to find out how big a buffer is needed for the generated string, *malloc()* a buffer to hold the  
37650 string, and then call *regerror()* again to get the string. Alternatively, it could allocate a fixed,  
37651 static buffer that is big enough to hold most strings, and then use *malloc()* to allocate a larger  
37652 buffer if it finds that this is too small.

37653 To match a pattern as described in the Shell and Utilities volume of IEEE Std 1003.1-200x, Section  
37654 2.13, Pattern Matching Notation, use the *fnmatch()* function.

#### 37655 RATIONALE

37656 The *regmatch()* function must fill in all *nmatch* elements of *pmatch*, where *nmatch* and *pmatch* are  
37657 supplied by the application, even if some elements of *pmatch* do not correspond to  
37658 subexpressions in *pattern*. The application writer should note that there is probably no reason  
37659 for using a value of *nmatch* that is larger than *preg->re\_nsub+1*.

37660 The REG\_NEWLINE flag supports a use of RE matching that is needed in some applications like  
37661 text editors. In such applications, the user supplies an RE asking the application to find a line  
37662 that matches the given expression. An anchor in such an RE anchors at the beginning or end of  
37663 any line. Such an application can pass a sequence of <newline>-separated lines to *regexec()* as a  
37664 single long string and specify REG\_NEWLINE to *regcomp()* to get the desired behavior. The  
37665 application must ensure that there are no explicit <newline>s in *pattern* if it wants to ensure that  
37666 any match occurs entirely within a single line.

37667 The REG\_NEWLINE flag affects the behavior of *regexec()*, but it is in the *cflags* parameter to  
37668 *regcomp()* to allow flexibility of implementation. Some implementations will want to generate  
37669 the same compiled RE in *regcomp()* regardless of the setting of REG\_NEWLINE and have  
37670 *regexec()* handle anchors differently based on the setting of the flag. Other implementations will  
37671 generate different compiled REs based on the REG\_NEWLINE.

37672 The REG\_ICASE flag supports the operations taken by the *grep -i* option and the historical  
37673 implementations of *ex* and *vi*. Including this flag will make it easier for application code to be  
37674 written that does the same thing as these utilities.

37675 The substrings reported in *pmatch[]* are defined using offsets from the start of the string rather  
37676 than pointers. Since this is a new interface, there should be no impact on historical  
37677 implementations or applications, and offsets should be just as easy to use as pointers. The  
37678 change to offsets was made to facilitate future extensions in which the string to be searched is  
37679 presented to *regexec()* in blocks, allowing a string to be searched that is not all in memory at  
37680 once.

37681 A new type **regoff\_t** is used for the elements of *pmatch[]* to ensure that the application can  
37682 represent either the largest possible array in memory (important for an application conforming  
37683 to the Shell and Utilities volume of IEEE Std 1003.1-200x) or the largest possible file (important  
37684 for an application using the extension where a file is searched in chunks).



37685 The standard developers rejected the inclusion of a *regsub()* function that would be used to do  
 37686 substitutions for a matched RE. While such a routine would be useful to some applications, its  
 37687 utility would be much more limited than the matching function described here. Both RE parsing  
 37688 and substitution are possible to implement without support other than that required by the  
 37689 ISO C standard, but matching is much more complex than substituting. The only difficult part of  
 37690 substitution, given the information supplied by *regexec()*, is finding the next character in a string  
 37691 when there can be multi-byte characters. That is a much larger issue, and one that needs a more  
 37692 general solution.

37693 The *errno* variable has not been used for error returns to avoid filling the *errno* name space for  
 37694 this feature.

37695 The interface is defined so that the matched substrings *rm\_sp* and *rm\_ep* are in a separate  
 37696 **regmatch\_t** structure instead of in **regex\_t**. This allows a single compiled RE to be used  
 37697 simultaneously in several contexts; in *main()* and a signal handler, perhaps, or in multiple  
 37698 threads of lightweight processes. (The *preg* argument to *regexec()* is declared with type **const**, so  
 37699 the implementation is not permitted to use the structure to store intermediate results.) It also  
 37700 allows an application to request an arbitrary number of substrings from an RE. The number of  
 37701 subexpressions in the RE is reported in *re\_nsub* in *preg*. With this change to *regexec()*,  
 37702 consideration was given to dropping the REG\_NOSUB flag since the user can now specify this  
 37703 with a zero *nmatch* argument to *regexec()*. However, keeping REG\_NOSUB allows an  
 37704 implementation to use a different (perhaps more efficient) algorithm if it knows in *regcomp()*  
 37705 that no subexpressions need be reported. The implementation is only required to fill in *pmatch* if  
 37706 *nmatch* is not zero and if REG\_NOSUB is not specified. Note that the **size\_t** type, as defined in  
 37707 the ISO C standard, is unsigned, so the description of *regexec()* does not need to address  
 37708 negative values of *nmatch*.

37709 REG\_NOTBOL was added to allow an application to do repeated searches for the same pattern  
 37710 in a line. If the pattern contains a circumflex character that should match the beginning of a line,  
 37711 then the pattern should only match when matched against the beginning of the line. Without  
 37712 the REG\_NOTBOL flag, the application could rewrite the expression for subsequent matches,  
 37713 but in the general case this would require parsing the expression. The need for REG\_NOTEOL is  
 37714 not as clear; it was added for symmetry.

37715 The addition of the *regerror()* function addresses the historical need for conforming application  
 37716 programs to have access to error information more than “Function failed to compile/match your  
 37717 RE for unknown reasons”.

37718 This interface provides for two different methods of dealing with error conditions. The specific  
 37719 error codes (REG\_EBRACE, for example), defined in **<regex.h>**, allow an application to recover  
 37720 from an error if it is so able. Many applications, especially those that use patterns supplied by a  
 37721 user, will not try to deal with specific error cases, but will just use *regerror()* to obtain a human-  
 37722 readable error message to present to the user.

37723 The *regerror()* function uses a scheme similar to *confstr()* to deal with the problem of allocating  
 37724 memory to hold the generated string. The scheme used by *strerror()* in the ISO C standard was  
 37725 considered unacceptable since it creates difficulties for multi-threaded applications.

37726 The *preg* argument is provided to *regerror()* to allow an implementation to generate a more  
 37727 descriptive message than would be possible with *errcode* alone. An implementation might, for  
 37728 example, save the character offset of the offending character of the pattern in a field of *preg*, and  
 37729 then include that in the generated message string. The implementation may also ignore *preg*.

37730 A REG\_FILENAME flag was considered, but omitted. This flag caused *regexec()* to match  
 37731 patterns as described in the Shell and Utilities volume of IEEE Std 1003.1-200x, Section 2.13,  
 37732 Pattern Matching Notation instead of REs. This service is now provided by the *fnmatch()*

- 37733 function.
- 37734 Notice that there is a difference in philosophy between the ISO POSIX-2:1993 standard and  
37735 IEEE Std 1003.1-200x in how to handle a bad regular expression. The ISO POSIX-2:1993 standard  
37736 says that many bad constructs produce undefined results, or that the interpretation is undefined.  
37737 IEEE Std 1003.1-200x, however, says that the interpretation of such REs is unspecified. The term  
37738 “undefined” means that the action by the application is an error, of similar severity to passing a  
37739 bad pointer to a function.
- 37740 The *regcomp()* and *regexec()* functions are required to accept any null-terminated string as the  
37741 *pattern* argument. If the meaning of the string is undefined, the behavior of the function is  
37742 unspecified. IEEE Std 1003.1-200x does not specify how the functions will interpret the pattern;  
37743 they might return error codes, or they might do pattern matching in some completely  
37744 unexpected way, but they should not do something like abort the process.
- 37745 **FUTURE DIRECTIONS**
- 37746 None.
- 37747 **SEE ALSO**
- 37748 *fnmatch()*, *glob()*, the Base Definitions volume of IEEE Std 1003.1-200x, `<regex.h>`, `<sys/types.h>`
- 37749 **CHANGE HISTORY**
- 37750 First released in Issue 4. Derived from the ISO POSIX-2 standard.
- 37751 **Issue 5**
- 37752 Moved from POSIX2 C-language Binding to BASE.
- 37753 **Issue 6**
- 37754 In the SYNOPSIS, the optional include of the `<sys/types.h>` header is removed.
- 37755 The following new requirements on POSIX implementations derive from alignment with the  
37756 Single UNIX Specification:
- 37757 • The requirement to include `<sys/types.h>` has been removed. Although `<sys/types.h>` was  
37758 required for conforming implementations of previous POSIX specifications, it was not  
37759 required for UNIX applications.
- 37760 The DESCRIPTION is updated to avoid use of the term “must” for application requirements.
- 37761 The REG\_ENOSYS constant is removed.
- 37762 The **restrict** keyword is added to the *regcomp()*, *regerror()*, and *regexec()* prototypes for  
37763 alignment with the ISO/IEC 9899:1999 standard.

37764 **NAME**

37765 remainder, remainderf, remainderl — remainder function

37766 **SYNOPSIS**

37767 #include &lt;math.h&gt;

37768 double remainder(double x, double y);

37769 float remainderf(float x, float y);

37770 long double remainderl(long double x, long double y);

37771 **DESCRIPTION**

37772 **CX** The functionality described on this reference page is aligned with the ISO C standard. Any  
 37773 conflict between the requirements described here and the ISO C standard is unintentional. This  
 37774 volume of IEEE Std 1003.1-200x defers to the ISO C standard.

37775 These functions shall return the floating-point remainder  $r=x-ny$  when  $y$  is non-zero. The value  
 37776  $n$  is the integral value nearest the exact value  $x/y$ . When  $|n-x/y| = 1/2$ , the value  $n$  is chosen to  
 37777 be even.

37778 The behavior of *remainder()* shall be independent of the rounding mode.37779 **RETURN VALUE**37780 Upon successful completion, these functions shall return the floating-point remainder  $r=x-ny$   
37781 when  $y$  is non-zero.37782 **MX** If  $x$  or  $y$  is NaN, a NaN shall be returned.37783 If  $x$  is infinite or  $y$  is 0 and the other is non-NaN, a domain error shall occur, and either a NaN (if  
37784 supported), or an implementation-defined value shall be returned.37785 **ERRORS**

37786 These functions shall fail if:

37787 <b>MX</b>	<b>Domain Error</b>	The $x$ argument is $\pm\text{Inf}$ , or the $y$ argument is $\pm 0$ and the other argument is non-NaN.
-----------------	---------------------	---

37789	37790	37791	37792	

37793 **EXAMPLES**

37794 None.

37795 **APPLICATION USAGE**37796 On error, the expressions (math\_errhandling & MATH\_ERRNO) and (math\_errhandling &  
37797 MATH\_ERREXCEPT) are independent of each other, but at least one of them must be non-zero.37798 **RATIONALE**

37799 None.

37800 **FUTURE DIRECTIONS**

37801 None.

37802 **SEE ALSO**37803 *abs()*, *div()*, *feclearexcept()*, *fetestexcept()*, *ldiv()*, the Base Definitions volume of |  
37804 IEEE Std 1003.1-200x, Section 4.18, Treatment of Error Conditions for Mathematical Functions, |  
37805 <math.h>

37806 **CHANGE HISTORY**

37807 First released in Issue 4, Version 2.

37808 **Issue 5**

37809 Moved from X/OPEN UNIX extension to BASE.

37810 **Issue 6**

37811 The *remainder()* function is no longer marked as an extension.

37812 The *remainderf()* and *remainderl()* functions are added for alignment with the ISO/IEC 9899:1999 standard.  
37813

37814 The DESCRIPTION, RETURN VALUE, ERRORS, and APPLICATION USAGE sections are  
37815 revised to align with the ISO/IEC 9899:1999 standard.

37816 IEC 60559:1989 standard floating-point extensions over the ISO/IEC 9899:1999 standard are  
37817 marked.

37818 **NAME**

37819 remove — remove a file

37820 **SYNOPSIS**

37821 #include &lt;stdio.h&gt;

37822 int remove(const char \*path);

37823 **DESCRIPTION**

37824 CX The functionality described on this reference page is aligned with the ISO C standard. Any  
 37825 conflict between the requirements described here and the ISO C standard is unintentional. This  
 37826 volume of IEEE Std 1003.1-200x defers to the ISO C standard.

37827 The *remove()* function shall cause the file named by the pathname pointed to by *path* to be no  
 37828 longer accessible by that name. A subsequent attempt to open that file using that name shall fail,  
 37829 unless it is created anew.

37830 CX If *path* does not name a directory, *remove(path)* shall be equivalent to *unlink(path)*.

37831 If *path* names a directory, *remove(path)* shall be equivalent to *rmdir(path)*.

37832 **RETURN VALUE**37833 CX Refer to *rmdir()* or *unlink()*.37834 **ERRORS**37835 CX Refer to *rmdir()* or *unlink()*.37836 **EXAMPLES**37837 **Removing Access to a File**37838 The following example shows how to remove access to a file named */home/cnd/old\_mods*.

37839 #include &lt;stdio.h&gt;

37840 int status;

37841 ...

37842 status = remove("/home/cnd/old\_mods");

37843 **APPLICATION USAGE**

37844 None.

37845 **RATIONALE**

37846 None.

37847 **FUTURE DIRECTIONS**

37848 None.

37849 **SEE ALSO**37850 *rmdir()*, *unlink()*, the Base Definitions volume of IEEE Std 1003.1-200x, <stdio.h>37851 **CHANGE HISTORY**

37852 First released in Issue 3.

37853 Entry included for alignment with the POSIX.1-1988 standard and the ISO C standard.

37854 **Issue 6**

37855 Extensions beyond the ISO C standard are now marked.

37856 The following new requirements on POSIX implementations derive from alignment with the  
 37857 Single UNIX Specification:

37858  
37859  
37860

- 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()*.

37861 **NAME**

37862       remque — remove an element from a queue

37863 **SYNOPSIS**

37864 xSI       #include &lt;search.h&gt;

37865       void remque(void \*element);

37866

37867 **DESCRIPTION**37868       Refer to *insque()*.

37869 **NAME**

37870 remquo, remquof, remquol — remainder functions

37871 **SYNOPSIS**

37872 #include &lt;math.h&gt;

37873 double remquo(double x, double y, int \*quo);

37874 float remquof(float x, float y, int \*quo);

37875 long double remquol(long double x, long double y, int \*quo);

37876 **DESCRIPTION**

37877 **CX** The functionality described on this reference page is aligned with the ISO C standard. Any  
 37878 conflict between the requirements described here and the ISO C standard is unintentional. This  
 37879 volume of IEEE Std 1003.1-200x defers to the ISO C standard.

37880 The *remquo()*, *remquof()*, and *remquol()* functions shall compute the same remainder as the  
 37881 *remainder()*, *remainderf()*, and *remainderl()* functions, respectively. In the object pointed to by  
 37882 *quo*, they store a value whose sign is the sign of  $x/y$  and whose magnitude is congruent modulo  
 37883  $2^n$  to the magnitude of the integral quotient of  $x/y$ , where  $n$  is an implementation-defined  
 37884 integer greater than or equal to 3.

37885 An application wishing to check for error situations should set *errno* to zero and call  
 37886 *feclearexcept*(FE\_ALL\_EXCEPT) before calling these functions. On return, if *errno* is non-zero or  
 37887 *fetestexcept*(FE\_INVALID | FE\_DIVBYZERO | FE\_OVERFLOW | FE\_UNDERFLOW) is non-  
 37888 zero, an error has occurred.

37889 **RETURN VALUE**37890 These functions shall return  $x \text{ REM } y$ .37891 **MX** If  $x$  or  $y$  is NaN, a NaN shall be returned.

37892 If  $x$  is  $\pm\text{Inf}$  or  $y$  is zero and the other argument is non-NaN, a domain error shall occur, and either  
 37893 a NaN (if supported), or an implementation-defined value shall be returned.

37894 **ERRORS**

37895 These functions shall fail if:

37896 **MX** Domain Error The  $x$  argument is  $\pm\text{Inf}$ , or the  $y$  argument is  $\pm 0$  and the other argument is  
 37897 non-NaN.

37898 If the integer expression (*math\_errhandling* & MATH\_ERRNO) is non-zero, |  
 37899 then *errno* shall be set to [EDOM]. If the integer expression (*math\_errhandling* |  
 37900 & MATH\_ERREXCEPT) is non-zero, then the invalid floating-point exception |  
 37901 shall be raised. |

37902 **EXAMPLES**

37903 None.

37904 **APPLICATION USAGE**

37905 On error, the expressions (*math\_errhandling* & MATH\_ERRNO) and (*math\_errhandling* &  
 37906 MATH\_ERREXCEPT) are independent of each other, but at least one of them must be non-zero.

37907 **RATIONALE**

37908 These functions are intended for implementing argument reductions which can exploit a few  
 37909 low-order bits of the quotient. Note that  $x$  may be so large in magnitude relative to  $y$  that an  
 37910 exact representation of the quotient is not practical.



37911 **FUTURE DIRECTIONS**

37912 None.

37913 **SEE ALSO**

37914 *feclearexcept()*, *fetetestexcept()*, *remainder()*, the Base Definitions volume of IEEE Std 1003.1-200x, |  
37915 Section 4.18, Treatment of Error Conditions for Mathematical Functions, <math.h> |

37916 **CHANGE HISTORY**

37917 First released in Issue 6. Derived from the ISO/IEC 9899:1999 standard.

## 37918 NAME

37919 rename — rename a file

## 37920 SYNOPSIS

37921 #include &lt;stdio.h&gt;

37922 int rename(const char \*old, const char \*new);

## 37923 DESCRIPTION

37924 cx The functionality described on this reference page is aligned with the ISO C standard. Any  
 37925 conflict between the requirements described here and the ISO C standard is unintentional. This  
 37926 volume of IEEE Std 1003.1-200x defers to the ISO C standard.

37927 The *rename()* function shall change the name of a file. The *old* argument points to the pathname  
 37928 of the file to be renamed. The *new* argument points to the new pathname of the file.

37929 cx If either the *old* or *new* argument names a symbolic link, *rename()* shall operate on the symbolic  
 37930 link itself, and shall not resolve the last component of the argument. If the *old* argument and the  
 37931 *new* argument resolve to the same existing file, *rename()* shall return successfully and perform no  
 37932 other action.

37933 If the *old* argument points to the pathname of a file that is not a directory, the *new* argument shall  
 37934 not point to the pathname of a directory. If the link named by the *new* argument exists, it shall be  
 37935 removed and *old* renamed to *new*. In this case, a link named *new* shall remain visible to other  
 37936 processes throughout the renaming operation and refer either to the file referred to by *new* or *old*  
 37937 before the operation began. Write access permission is required for both the directory containing  
 37938 *old* and the directory containing *new*.

37939 If the *old* argument points to the pathname of a directory, the *new* argument shall not point to the  
 37940 pathname of a file that is not a directory. If the directory named by the *new* argument exists, it  
 37941 shall be removed and *old* renamed to *new*. In this case, a link named *new* shall exist throughout  
 37942 the renaming operation and shall refer either to the directory referred to by *new* or *old* before the  
 37943 operation began. If *new* names an existing directory, it shall be required to be an empty directory.

37944 If the *old* argument points to a pathname of a symbolic link, the symbolic link shall be renamed.  
 37945 If the *new* argument points to a pathname of a symbolic link, the symbolic link shall be removed.

37946 The *new* pathname shall not contain a path prefix that names *old*. Write access permission is  
 37947 required for the directory containing *old* and the directory containing *new*. If the *old* argument  
 37948 points to the pathname of a directory, write access permission may be required for the directory  
 37949 named by *old*, and, if it exists, the directory named by *new*.

37950 If the link named by the *new* argument exists and the file's link count becomes 0 when it is  
 37951 removed and no process has the file open, the space occupied by the file shall be freed and the  
 37952 file shall no longer be accessible. If one or more processes have the file open when the last link is  
 37953 removed, the link shall be removed before *rename()* returns, but the removal of the file contents  
 37954 shall be postponed until all references to the file are closed.

37955 Upon successful completion, *rename()* shall mark for update the *st\_ctime* and *st\_mtime* fields of  
 37956 the parent directory of each file.

37957 If the *rename()* function fails for any reason other than [EIO], any file named by *new* shall be  
 37958 unaffected.

## 37959 RETURN VALUE

37960 cx Upon successful completion, *rename()* shall return 0; otherwise,  $-1$  shall be returned, *errno* shall  
 37961 be set to indicate the error, and neither the file named by *old* nor the file named by *new* shall be  
 37962 changed or created.

37963 **ERRORS**37964 The *rename()* function shall fail if:

37965 CX [EACCES] A component of either path prefix denies search permission; or one of the  
 37966 directories containing *old* or *new* denies write permissions; or, write  
 37967 permission is required and is denied for a directory pointed to by the *old* or  
 37968 *new* arguments.

37969 CX [EBUSY] The directory named by *old* or *new* is currently in use by the system or another  
 37970 process, and the implementation considers this an error.

37971 CX [EEXIST] or [ENOTEMPTY]  
 37972 The link named by *new* is a directory that is not an empty directory.

37973 CX [EINVAL] The *new* directory pathname contains a path prefix that names the *old* |  
 37974 directory.

37975 CX [EIO] A physical I/O error has occurred.

37976 CX [EISDIR] The *new* argument points to a directory and the *old* argument points to a file  
 37977 that is not a directory.

37978 CX [ELOOP] A loop exists in symbolic links encountered during resolution of the *path*  
 37979 argument.

37980 CX [EMLINK] The file named by *old* is a directory, and the link count of the parent directory  
 37981 of *new* would exceed {LINK\_MAX}.

37982 CX [ENAMETOOLONG]  
 37983 The length of the *old* or *new* argument exceeds {PATH\_MAX} or a pathname |  
 37984 component is longer than {NAME\_MAX}.

37985 CX [ENOENT] The link named by *old* does not name an existing file, or either *old* or *new*  
 37986 points to an empty string.

37987 CX [ENOSPC] The directory that would contain *new* cannot be extended.

37988 CX [ENOTDIR] A component of either path prefix is not a directory; or the *old* argument  
 37989 names a directory and *new* argument names a non-directory file.

37990 XSI [EPERM] or [EACCES]  
 37991 The S\_ISVTX flag is set on the directory containing the file referred to by *old*  
 37992 and the caller is not the file owner, nor is the caller the directory owner, nor  
 37993 does the caller have appropriate privileges; or *new* refers to an existing file, the  
 37994 S\_ISVTX flag is set on the directory containing this file, and the caller is not  
 37995 the file owner, nor is the caller the directory owner, nor does the caller have  
 37996 appropriate privileges.

37997 CX [EROFS] The requested operation requires writing in a directory on a read-only file  
 37998 system.

37999 CX [EXDEV] The links named by *new* and *old* are on different file systems and the  
 38000 implementation does not support links between file systems.

38001 The *rename()* function may fail if:

38002 XSI [EBUSY] The file named by the *old* or *new* arguments is a named STREAM.

38003 CX [ELOOP] More than {SYMLOOP\_MAX} symbolic links were encountered during  
 38004 resolution of the *path* argument.

38005 CX [ENAMETOOLONG]  
 38006 As a result of encountering a symbolic link in resolution of the *path* argument,  
 38007 the length of the substituted pathname string exceeded {PATH\_MAX}. |

38008 CX [ETXTBSY] The file to be renamed is a pure procedure (shared text) file that is being  
 38009 executed. |

38010 **EXAMPLES**38011 **Renaming a File**

38012 The following example shows how to rename a file named `/home/cnd/mod1` to  
 38013 `/home/cnd/mod2`.

```
38014 #include <stdio.h>
38015 int status;
38016 ...
38017 status = rename("/home/cnd/mod1", "/home/cnd/mod2");
```

38018 **APPLICATION USAGE**

38019 Some implementations mark for update the *st\_ctime* field of renamed files and some do not. |  
 38020 Applications which make use of the *st\_ctime* field may behave differently with respect to |  
 38021 renamed files unless they are designed to allow for either behavior. |

38022 **RATIONALE**

38023 This *rename()* function is equivalent for regular files to that defined by the ISO C standard. Its  
 38024 inclusion here expands that definition to include actions on directories and specifies behavior  
 38025 when the *new* parameter names a file that already exists. That specification requires that the  
 38026 action of the function be atomic.

38027 One of the reasons for introducing this function was to have a means of renaming directories  
 38028 while permitting implementations to prohibit the use of *link()* and *unlink()* with directories,  
 38029 thus constraining links to directories to those made by *mkdir()*.

38030 The specification that if *old* and *new* refer to the same file is intended to guarantee that:

```
38031 rename("x", "x");
```

38032 does not remove the file.

38033 Renaming dot or dot-dot is prohibited in order to prevent cyclical file system paths.

38034 See also the descriptions of [ENOTEMPTY] and [ENAMETOOLONG] in *rmdir()* and [EBUSY] in  
 38035 *unlink()*. For a discussion of [EXDEV], see *link()*.

38036 **FUTURE DIRECTIONS**

38037 None.

38038 **SEE ALSO**

38039 *link()*, *rmdir()*, *symlink()*, *unlink()*, the Base Definitions volume of IEEE Std 1003.1-200x,  
 38040 `<stdio.h>`

38041 **CHANGE HISTORY**

38042 First released in Issue 3.

38043 Entry included for alignment with the POSIX.1-1988 standard.

38044 **Issue 5**

38045 The [EBUSY] error is added to the “may fail” part of the ERRORS section.

38046 **Issue 6**

38047 Extensions beyond the ISO C standard are now marked.

38048 The following new requirements on POSIX implementations derive from alignment with the |  
38049 Single UNIX Specification:

- 38050 • The [EIO] mandatory error condition is added.
- 38051 • The [ELOOP] mandatory error condition is added.
- 38052 • A second [ENAMETOOLONG] is added as an optional error condition.
- 38053 • The [ETXTBSY] optional error condition is added.

38054 The following changes were made to align with the IEEE P1003.1a draft standard:

- 38055 • Details are added regarding the treatment of symbolic links.
- 38056 • The [ELOOP] optional error condition is added.

38057 The DESCRIPTION is updated to avoid use of the term “must” for application requirements.

38058 **NAME**

38059           rewind — reset file position indicator in a stream

38060 **SYNOPSIS**

38061           #include &lt;stdio.h&gt;

38062           void rewind(FILE \*stream);

38063 **DESCRIPTION**

38064 cx       The functionality described on this reference page is aligned with the ISO C standard. Any  
38065       conflict between the requirements described here and the ISO C standard is unintentional. This  
38066       volume of IEEE Std 1003.1-200x defers to the ISO C standard.

38067       The call:

38068       rewind(stream)

38069       shall be equivalent to:

38070       (void) fseek(stream, 0L, SEEK\_SET)

38071       except that *rewind()* shall also clear the error indicator. |

38072 cx       Since *rewind()* does not return a value, an application wishing to detect errors should clear *errno*, |  
38073       then call *rewind()*, and if *errno* is non-zero, assume an error has occurred.

38074 **RETURN VALUE**38075       The *rewind()* function shall not return a value.38076 **ERRORS**38077 cx       Refer to *fseek()* with the exception of [EINVAL] which does not apply.38078 **EXAMPLES**

38079       None.

38080 **APPLICATION USAGE**

38081       None.

38082 **RATIONALE**

38083       None.

38084 **FUTURE DIRECTIONS**

38085       None.

38086 **SEE ALSO**38087       *fseek()*, the Base Definitions volume of IEEE Std 1003.1-200x, <stdio.h>38088 **CHANGE HISTORY**

38089       First released in Issue 1. Derived from Issue 1 of the SVID.

38090 **Issue 6**

38091       Extensions beyond the ISO C standard are now marked.

38092 **NAME**

38093           rewinddir — reset position of directory stream to the beginning of a directory

38094 **SYNOPSIS**

38095           #include &lt;dirent.h&gt;

38096           void rewinddir(DIR \*dirp);

38097 **DESCRIPTION**

38098           The *rewinddir()* function shall reset the position of the directory stream to which *dirp* refers to the beginning of the directory. It shall also cause the directory stream to refer to the current state of the corresponding directory, as a call to *opendir()* would have done. If *dirp* does not refer to a directory stream, the effect is undefined.

38102           After a call to the *fork()* function, either the parent or child (but not both) may continue processing the directory stream using *readdir()*, *rewinddir()*, or *seekdir()*. If both the parent and child processes use these functions, the result is undefined.

38105 **RETURN VALUE**38106           The *rewinddir()* function shall not return a value.38107 **ERRORS**

38108           No errors are defined.

38109 **EXAMPLES**

38110           None.

38111 **APPLICATION USAGE**

38112           The *rewinddir()* function should be used in conjunction with *opendir()*, *readdir()*, and *closedir()* to examine the contents of the directory. This method is recommended for portability.

38114 **RATIONALE**

38115           None.

38116 **FUTURE DIRECTIONS**

38117           None.

38118 **SEE ALSO**

38119           *closedir()*, *opendir()*, *readdir()*, the Base Definitions volume of IEEE Std 1003.1-200x, <dirent.h>  
38120           <sys/types.h>

38121 **CHANGE HISTORY**

38122           First released in Issue 2.

38123 **Issue 6**

38124           In the SYNOPSIS, the optional include of the &lt;sys/types.h&gt; header is removed.

38125           The following new requirements on POSIX implementations derive from alignment with the Single UNIX Specification:

- 38127           • 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.

38130 **NAME**38131 rindex — character string operations (**LEGACY**)38132 **SYNOPSIS**38133 XSI `#include <strings.h>`38134 `char *rindex(const char *s, int c);`

38135

38136 **DESCRIPTION**38137 The *rindex()* function shall be equivalent to *strchr()*.38138 **RETURN VALUE**38139 Refer to *strchr()*.38140 **ERRORS**38141 Refer to *strchr()*.38142 **EXAMPLES**

38143 None.

38144 **APPLICATION USAGE**38145 *strchr()* is preferred over this function.38146 For maximum portability, it is recommended to replace the function call to *rindex()* as follows:38147 `#define rindex(a,b) strchr((a),(b))`38148 **RATIONALE**

38149 None.

38150 **FUTURE DIRECTIONS**

38151 This function may be withdrawn in a future version.

38152 **SEE ALSO**38153 *strchr()*, the Base Definitions volume of IEEE Std 1003.1-200x, `<strings.h>`38154 **CHANGE HISTORY**

38155 First released in Issue 4, Version 2.

38156 **Issue 5**

38157 Moved from X/OPEN UNIX extension to BASE.

38158 **Issue 6**

38159 This function is marked LEGACY.



38160 **NAME**

38161 rint, rintf, rintl — round-to-nearest integral value

38162 **SYNOPSIS**

38163 #include &lt;math.h&gt;

38164 double rint(double x);

38165 float rintf(float x);

38166 long double rintl(long double x);

38167 **DESCRIPTION**

38168 **CX** The functionality described on this reference page is aligned with the ISO C standard. Any  
 38169 conflict between the requirements described here and the ISO C standard is unintentional. This  
 38170 volume of IEEE Std 1003.1-200x defers to the ISO C standard.

38171 These functions shall return the integral value (represented as a **double**) nearest  $x$  in the  
 38172 direction of the current rounding mode. The current rounding mode is implementation-defined.

38173 If the current rounding mode rounds toward negative infinity, then *rint()* shall be equivalent to  
 38174 *floor()*. If the current rounding mode rounds toward positive infinity, then *rint()* shall be  
 38175 equivalent to *ceil()*.

38176 These functions differ from the *nearbyint()*, *nearbyintf()*, and *nearbyintl()* functions only in that  
 38177 they may raise the inexact floating-point exception if the result differs in value from the  
 38178 argument.

38179 An application wishing to check for error situations should set *errno* to zero and call  
 38180 *feclearexcept*(FE\_ALL\_EXCEPT) before calling these functions. On return, if *errno* is non-zero or  
 38181 *fetestexcept*(FE\_INVALID | FE\_DIVBYZERO | FE\_OVERFLOW | FE\_UNDERFLOW) is non-  
 38182 zero, an error has occurred.

38183 **RETURN VALUE**

38184 Upon successful completion, these functions shall return the integer (represented as a double  
 38185 precision number) nearest  $x$  in the direction of the current rounding mode.

38186 **MX** If  $x$  is NaN, a NaN shall be returned.

38187 If  $x$  is  $\pm 0$ , or  $\pm \text{Inf}$ ,  $x$  shall be returned.

38188 **XSI** If the correct value would cause overflow, a range error shall occur and *rint()*, *rintf()*, and *rintl()*  
 38189 shall return the value of the macro HUGE\_VAL, HUGE\_VALF, and HUGE\_VALL, respectively.

38190 **ERRORS**

38191 These functions shall fail if:

38192 **XSI** **Range Error** The result would cause an overflow.

38193 If the integer expression (math\_errhandling & MATH\_ERRNO) is non-zero,  
 38194 then *errno* shall be set to [ERANGE]. If the integer expression  
 38195 (math\_errhandling & MATH\_ERREXCEPT) is non-zero, then the overflow  
 38196 floating-point exception shall be raised.

38197 **EXAMPLES**

38198 None.

38199 **APPLICATION USAGE**

38200 On error, the expressions `(math_errhandling & MATH_ERRNO)` and `(math_errhandling &`  
38201 `MATH_ERREXCEPT)` are independent of each other, but at least one of them must be non-zero.

38202 **RATIONALE**

38203 None.

38204 **FUTURE DIRECTIONS**

38205 None.

38206 **SEE ALSO**

38207 *abs()*, *ceil()*, *feclearexcept()*, *fetestexcept()*, *nearbyint()*, *floor()*, *isnan()*, the Base Definitions volume |  
38208 of IEEE Std 1003.1-200x, Section 4.18, Treatment of Error Conditions for Mathematical Functions, |  
38209 `<math.h>`

38210 **CHANGE HISTORY**

38211 First released in Issue 4, Version 2.

38212 **Issue 5**

38213 Moved from X/OPEN UNIX extension to BASE.

38214 **Issue 6**

38215 The following changes are made for alignment with the ISO/IEC 9899:1999 standard:

- 38216 • The *rintf()* and *rintl()* functions are added.
- 38217 • The *rint()* function is no longer marked as an extension.
- 38218 • The DESCRIPTION, RETURN VALUE, ERRORS, and APPLICATION USAGE sections are  
38219 revised to align with the ISO/IEC 9899:1999 standard.
- 38220 IEC 60559:1989 standard floating-point extensions over the ISO/IEC 9899:1999 standard are  
38221 marked.

38222 **NAME**

38223 rmdir — remove a directory

38224 **SYNOPSIS**

38225 #include &lt;unistd.h&gt;

38226 int rmdir(const char \*path);

38227 **DESCRIPTION**

38228 The *rmdir()* function shall remove a directory whose name is given by *path*. The directory shall  
 38229 be removed only if it is an empty directory. |

38230 If the directory is the root directory or the current working directory of any process, it is  
 38231 unspecified whether the function succeeds, or whether it shall fail and set *errno* to [EBUSY].

38232 If *path* names a symbolic link, then *rmdir()* shall fail and set *errno* to [ENOTDIR].

38233 If the *path* argument refers to a path whose final component is either dot or dot-dot, *rmdir()* shall  
 38234 fail.

38235 If the directory's link count becomes 0 and no process has the directory open, the space occupied  
 38236 by the directory shall be freed and the directory shall no longer be accessible. If one or more  
 38237 processes have the directory open when the last link is removed, the dot and dot-dot entries, if  
 38238 present, shall be removed before *rmdir()* returns and no new entries may be created in the  
 38239 directory, but the directory shall not be removed until all references to the directory are closed.

38240 If the directory is not an empty directory, *rmdir()* shall fail and set *errno* to [EEXIST] or  
 38241 [ENOTEMPTY].

38242 Upon successful completion, the *rmdir()* function shall mark for update the *st\_ctime* and  
 38243 *st\_mtime* fields of the parent directory.

38244 **RETURN VALUE**

38245 Upon successful completion, the function *rmdir()* shall return 0. Otherwise, -1 shall be returned,  
 38246 and *errno* set to indicate the error. If -1 is returned, the named directory shall not be changed.

38247 **ERRORS**

38248 The *rmdir()* function shall fail if:

38249 [EACCES] Search permission is denied on a component of the path prefix, or write  
 38250 permission is denied on the parent directory of the directory to be removed.

38251 [EBUSY] The directory to be removed is currently in use by the system or some process  
 38252 and the implementation considers this to be an error.

38253 [EEXIST] or [ENOTEMPTY]  
 38254 The *path* argument names a directory that is not an empty directory, or there  
 38255 are hard links to the directory other than dot or a single entry in dot-dot.

38256 [EINVAL] The *path* argument contains a last component that is dot.

38257 [EIO] A physical I/O error has occurred.

38258 [ELOOP] A loop exists in symbolic links encountered during resolution of the *path*  
 38259 argument.

38260 [ENAMETOOLONG]  
 38261 The length of the *path* argument exceeds {PATH\_MAX} or a pathname  
 38262 component is longer than |

38263 NAME\_MAX

38264	[ENOENT]	A component of <i>path</i> does not name an existing file, or the <i>path</i> argument names a nonexistent directory or points to an empty string.
38265		
38266	[ENOTDIR]	A component of <i>path</i> is not a directory.
38267 XSI	[EPERM] or [EACCES]	
38268		The S_ISVTX flag is set on the parent directory of the directory to be removed and the caller is not the owner of the directory to be removed, nor is the caller the owner of the parent directory, nor does the caller have the appropriate privileges.
38269		
38270		
38271		
38272	[EROFS]	The directory entry to be removed resides on a read-only file system.
38273		The <i>rmdir()</i> function may fail if:
38274	[ELOOP]	More than {SYMLOOP_MAX} symbolic links were encountered during resolution of the <i>path</i> argument.
38275		
38276	[ENAMETOOLONG]	
38277		As a result of encountering a symbolic link in resolution of the <i>path</i> argument,
38278		the length of the substituted pathname string exceeded {PATH_MAX}.

### 38279 EXAMPLES

#### 38280 Removing a Directory

38281 The following example shows how to remove a directory named `/home/cnd/mod1`.

```
38282 #include <unistd.h>
38283
38284 int status;
38285 ...
38286 status = rmdir("/home/cnd/mod1");
```

### 38286 APPLICATION USAGE

38287 None.

### 38288 RATIONALE

38289 The *rmdir()* and *rename()* functions originated in 4.2 BSD, and they used [ENOTEMPTY] for the condition when the directory to be removed does not exist or *new* already exists. When the 1984 /usr/group standard was published, it contained [EEXIST] instead. When these functions were adopted into System V, the 1984 /usr/group standard was used as a reference. Therefore, several existing applications and implementations support/use both forms, and no agreement could be reached on either value. All implementations are required to supply both [EEXIST] and [ENOTEMPTY] in `<errno.h>` with distinct values, so that applications can use both values in C-language `case` statements.

38297 The meaning of deleting *pathname/dot* is unclear, because the name of the file (directory) in the parent directory to be removed is not clear, particularly in the presence of multiple links to a directory.

38300 IEEE Std 1003.1-200x was silent with regard to the behavior of *rmdir()* when there are multiple hard links to the directory being removed. The requirement to set *errno* to [EEXIST] or [ENOTEMPTY] clarifies the behavior in this case.

38303 If the process' current working directory is being removed, that should be an allowed error.

38304 Virtually all existing implementations detect [ENOTEMPTY] or the case of dot-dot. The text in Section 2.3 (on page 471) about returning any one of the possible errors permits that behavior to continue. The [ELOOP] error may be returned if more than {SYMLOOP\_MAX} symbolic links

38307 are encountered during resolution of the *path* argument.

38308 **FUTURE DIRECTIONS**

38309 None.

38310 **SEE ALSO**

38311 *mkdir()*, *remove()*, *unlink()*, the Base Definitions volume of IEEE Std 1003.1-200x, <**unistd.h**>

38312 **CHANGE HISTORY**

38313 First released in Issue 3.

38314 Entry included for alignment with the POSIX.1-1988 standard.

38315 **Issue 6**

38316 The following new requirements on POSIX implementations derive from alignment with the |  
38317 Single UNIX Specification:

- 38318 • The DESCRIPTION is updated to indicate the results of naming a symbolic link in *path*.
- 38319 • The [EIO] mandatory error condition is added.
- 38320 • The [ELOOP] mandatory error condition is added.
- 38321 • A second [ENAMETOOLONG] is added as an optional error condition.

38322 The following changes were made to align with the IEEE P1003.1a draft standard:

- 38323 • The [ELOOP] optional error condition is added.

38324 **NAME**

38325 round, roundf, roundl — round to nearest integer value in floating-point format

38326 **SYNOPSIS**

38327 #include &lt;math.h&gt;

38328 double round(double x);

38329 float roundf(float x);

38330 long double roundl(long double x);

38331 **DESCRIPTION**38332 **CX** The functionality described on this reference page is aligned with the ISO C standard. Any  
38333 conflict between the requirements described here and the ISO C standard is unintentional. This  
38334 volume of IEEE Std 1003.1-200x defers to the ISO C standard.38335 These functions shall round their argument to the nearest integer value in floating-point format,  
38336 rounding halfway cases away from zero, regardless of the current rounding direction.38337 An application wishing to check for error situations should set *errno* to zero and call  
38338 *feclearexcept*(FE\_ALL\_EXCEPT) before calling these functions. On return, if *errno* is non-zero or  
38339 *fetestexcept*(FE\_INVALID | FE\_DIVBYZERO | FE\_OVERFLOW | FE\_UNDERFLOW) is non-  
38340 zero, an error has occurred.38341 **RETURN VALUE**

38342 Upon successful completion, these functions shall return the rounded integer value.

38343 **MX** If *x* is NaN, a NaN shall be returned.38344 If *x* is  $\pm 0$ , or  $\pm \text{Inf}$ , *x* shall be returned.38345 **XSI** If the correct value would cause overflow, a range error shall occur and *round()*, *roundf()*, and  
38346 *roundl()* shall return the value of the macro HUGE\_VAL, HUGE\_VALF, and HUGE\_VALL,  
38347 respectively.38348 **ERRORS**

38349 These functions may fail if:

38350 **XSI** **Range Error** The result overflows.38351 If the integer expression (math\_errhandling & MATH\_ERRNO) is non-zero, |  
38352 then *errno* shall be set to [ERANGE]. If the integer expression |  
38353 (math\_errhandling & MATH\_ERREXCEPT) is non-zero, then the overflow |  
38354 floating-point exception shall be raised. |38355 **EXAMPLES**

38356 None.

38357 **APPLICATION USAGE**38358 On error, the expressions (math\_errhandling & MATH\_ERRNO) and (math\_errhandling &  
38359 MATH\_ERREXCEPT) are independent of each other, but at least one of them must be non-zero.38360 **RATIONALE**

38361 None.

38362 **FUTURE DIRECTIONS**

38363 None.

38364 **SEE ALSO**

38365 *feclearexcept()*, *fetestexcept()*, the Base Definitions volume of IEEE Std 1003.1-200x, Section 4.18, |  
38366 Treatment of Error Conditions for Mathematical Functions, <math.h> |

38367 **CHANGE HISTORY**

38368 First released in Issue 6. Derived from the ISO/IEC 9899:1999 standard.

## 38369 NAME

38370 scalb — load exponent of a radix-independent floating-point number

## 38371 SYNOPSIS

38372 OB XSI `#include <math.h>`38373 `double scalb(double x, double n);`

38374

## 38375 DESCRIPTION

38376 The *scalb()* function shall compute  $x \cdot r^n$ , where  $r$  is the radix of the machine's floating-point  
 38377 arithmetic. When  $r$  is 2, *scalb()* shall be equivalent to *ldexp()*. The value of  $r$  is FLT\_RADIX  
 38378 which is defined in `<float.h>`.

38379 An application wishing to check for error situations should set *errno* to zero and call  
 38380 *feclearexcept*(FE\_ALL\_EXCEPT) before calling these functions. On return, if *errno* is non-zero or  
 38381 *fetestexcept*(FE\_INVALID | FE\_DIVBYZERO | FE\_OVERFLOW | FE\_UNDERFLOW) is non-  
 38382 zero, an error has occurred.

## 38383 RETURN VALUE

38384 Upon successful completion, the *scalb()* function shall return  $x \cdot r^n$ .38385 If  $x$  or  $n$  is NaN, a NaN shall be returned.38386 If  $n$  is zero,  $x$  shall be returned.38387 If  $x$  is  $\pm\text{Inf}$  and  $n$  is not  $-\text{Inf}$ ,  $x$  shall be returned.38388 If  $x$  is  $\pm 0$  and  $n$  is not  $+\text{Inf}$ ,  $x$  shall be returned.

38389 If  $x$  is  $\pm 0$  and  $n$  is  $+\text{Inf}$ , a domain error shall occur, and either a NaN (if supported), or an  
 38390 implementation-defined value shall be returned.

38391 If  $x$  is  $\pm\text{Inf}$  and  $n$  is  $-\text{Inf}$ , a domain error shall occur, and either a NaN (if supported), or an  
 38392 implementation-defined value shall be returned.

38393 If the result would cause an overflow, a range error shall occur and  $\pm\text{HUGE\_VAL}$  (according to  
 38394 the sign of  $x$ ) shall be returned.

38395 If the correct value would cause underflow, and is representable, a range error may occur and  
 38396 the correct value shall be returned.

38397 If the correct value would cause underflow, and is not representable, a range error may occur,  
 38398 and 0.0 shall be returned.

## 38399 ERRORS

38400 The *scalb()* function shall fail if:38401 Domain Error If  $x$  is zero and  $n$  is  $+\text{Inf}$ , or  $x$  is  $\text{Inf}$  and  $n$  is  $-\text{Inf}$ .

38402 If the integer expression (*math\_errhandling* & MATH\_ERRNO) is non-zero, |  
 38403 then *errno* shall be set to [EDOM]. If the integer expression (*math\_errhandling* |  
 38404 & MATH\_ERREXCEPT) is non-zero, then the invalid floating-point exception |  
 38405 shall be raised. |

38406 Range Error The result would overflow.

38407 If the integer expression (*math\_errhandling* & MATH\_ERRNO) is non-zero, |  
 38408 then *errno* shall be set to [ERANGE]. If the integer expression |  
 38409 (*math\_errhandling* & MATH\_ERREXCEPT) is non-zero, then the overflow |  
 38410 floating-point exception shall be raised. |



38411 The *scalb()* function may fail if:

38412 Range Error The result underflows.

38413 If the integer expression (*math\_errhandling* & MATH\_ERRNO) is non-zero, |  
 38414 then *errno* shall be set to [ERANGE]. If the integer expression |  
 38415 (*math\_errhandling* & MATH\_ERREXCEPT) is non-zero, then the underflow |  
 38416 floating-point exception shall be raised. |

38417 **EXAMPLES**

38418 None.

38419 **APPLICATION USAGE**

38420 Applications should use either *scalbln()*, *scalblnf()*, or *scalblnl()* in preference to this function.

38421 IEEE Std 1003.1-200x only defines the behavior for the *scalb()* function when the *n* argument is  
 38422 an integer, a NaN, or Inf. The behavior of other values for the *n* argument is unspecified.

38423 On error, the expressions (*math\_errhandling* & MATH\_ERRNO) and (*math\_errhandling* &  
 38424 MATH\_ERREXCEPT) are independent of each other, but at least one of them must be non-zero.

38425 **RATIONALE**

38426 None.

38427 **FUTURE DIRECTIONS**

38428 None.

38429 **SEE ALSO**

38430 *feclearexcept()*, *fetetestexcept()*, *ilogb()*, *ldexp()*, *logb()*, *scalbln()*, the Base Definitions volume of |  
 38431 IEEE Std 1003.1-200x, Section 4.18, Treatment of Error Conditions for Mathematical Functions, |  
 38432 <float.h>, <math.h>

38433 **CHANGE HISTORY**

38434 First released in Issue 4, Version 2.

38435 **Issue 5**

38436 Moved from X/OPEN UNIX extension to BASE.

38437 The DESCRIPTION is updated to indicate how an application should check for an error. This  
 38438 text was previously published in the APPLICATION USAGE section.

38439 **Issue 6**

38440 This function is marked obsolescent.

38441 Although this function is not part of the ISO/IEC 9899:1999 standard, the RETURN VALUE and  
 38442 ERROR sections are updated to align with the error handling in ISO/IEC 9899:1999 standard.

38443 **NAME**

38444 scalbn, scalblnf, scalblnl, scalbn, scalbnf, scalbnl, — compute exponent using FLT\_RADIX

38445 **SYNOPSIS**

38446 #include &lt;math.h&gt;

38447 double scalbn(double x, long n);

38448 float scalblnf(float x, long n);

38449 long double scalblnl(long double x, long n);

38450 double scalbn(double x, int n);

38451 float scalbnf(float x, int n);

38452 long double scalbnl(long double x, int n);

38453 **DESCRIPTION**38454 **CX** The functionality described on this reference page is aligned with the ISO C standard. Any  
38455 conflict between the requirements described here and the ISO C standard is unintentional. This  
38456 volume of IEEE Std 1003.1-200x defers to the ISO C standard.38457 These functions shall compute  $x * FLT\_RADIX^n$  efficiently, not normally by computing  
38458  $FLT\_RADIX^n$  explicitly.38459 An application wishing to check for error situations should set *errno* to zero and call  
38460 *feclearexcept*(FE\_ALL\_EXCEPT) before calling these functions. On return, if *errno* is non-zero or  
38461 *fetestexcept*(FE\_INVALID | FE\_DIVBYZERO | FE\_OVERFLOW | FE\_UNDERFLOW) is non-  
38462 zero, an error has occurred.38463 **RETURN VALUE**38464 Upon successful completion, these functions shall return  $x * FLT\_RADIX^n$ .38465 If the result would cause overflow, a range error shall occur and these functions shall return  
38466  $\pm HUGUE\_VAL$ ,  $\pm HUGUE\_VALF$ , and  $\pm HUGUE\_VALL$  (according to the sign of *x*) as appropriate for  
38467 the return type of the function.38468 If the correct value would cause underflow, and is not representable, a range error may occur,  
38469 **MX** and either 0.0 (if supported), or an implementation-defined value shall be returned.38470 **MX** If *x* is NaN, a NaN shall be returned.38471 If *x* is  $\pm 0$ , or  $\pm Inf$ , *x* shall be returned.38472 If *n* is 0, *x* shall be returned.38473 If the correct value would cause underflow, and is representable, a range error may occur and  
38474 the correct value shall be returned.38475 **ERRORS**

38476 These functions shall fail if:

38477 Range Error The result overflows.

38478 If the integer expression (math\_errhandling & MATH\_ERRNO) is non-zero, |  
38479 then *errno* shall be set to [ERANGE]. If the integer expression |  
38480 (math\_errhandling & MATH\_ERREXCEPT) is non-zero, then the overflow |  
38481 floating-point exception shall be raised. |

38482 These functions may fail if:

38483 Range Error The result underflows.

38484 If the integer expression (math\_errhandling & MATH\_ERRNO) is non-zero, |  
38485 then *errno* shall be set to [ERANGE]. If the integer expression |

38486 (math\_errhandling & MATH\_ERREXCEPT) is non-zero, then the underflow |  
38487 floating-point exception shall be raised. |

#### 38488 EXAMPLES

38489 None.

#### 38490 APPLICATION USAGE

38491 On error, the expressions (math\_errhandling & MATH\_ERRNO) and (math\_errhandling &  
38492 MATH\_ERREXCEPT) are independent of each other, but at least one of them must be non-zero.

#### 38493 RATIONALE

38494 These functions are named so as to avoid conflicting with the historical definition of the *scalb()*  
38495 function from the Single UNIX Specification. The difference is that the *scalb()* function has  
38496 second argument of **double** instead of **int**. The *scalb()* function is not part of ISO C standard. |  
38497 The three functions whose second type is **long** are provided because the factor required to scale |  
38498 from the smallest positive floating-point value to the largest finite one, on many |  
38499 implementations, is too large to represent in the minimum-width **int** format.

#### 38500 FUTURE DIRECTIONS

38501 None.

#### 38502 SEE ALSO

38503 *feclearexcept()*, *fetestexcept()*, *scalb()*, the Base Definitions volume of IEEE Std 1003.1-200x, Section |  
38504 4.18, Treatment of Error Conditions for Mathematical Functions, <math.h> |

#### 38505 CHANGE HISTORY

38506 First released in Issue 6. Derived from the ISO/IEC 9899:1999 standard.

38507 **NAME**

38508           scalbn, scalbnf, scalbnl, — compute exponent using FLT\_RADIX

38509 **SYNOPSIS**

38510           #include &lt;math.h&gt;

38511           double scalbn(double *x*, int *n*);38512           float scalbnf(float *x*, int *n*);38513           long double scalbnl(long double *x*, int *n*);38514 **DESCRIPTION**38515           Refer to *scalbln()*.

38516 **NAME**

38517       scanf — convert formatted input

38518 **SYNOPSIS**

38519       #include &lt;stdio.h&gt;

38520       int scanf(const char \*restrict *format*, ... );38521 **DESCRIPTION**38522       Refer to *fscanf()*.

38523 **NAME**

38524 sched\_get\_priority\_max, sched\_get\_priority\_min — get priority limits (**REALTIME**)

38525 **SYNOPSIS**

```
38526 PS #include <sched.h>
```

```
38527 int sched_get_priority_max(int policy);
```

```
38528 int sched_get_priority_min(int policy);
```

```
38529
```

38530 **DESCRIPTION**

38531 The *sched\_get\_priority\_max()* and *sched\_get\_priority\_min()* functions shall return the appropriate  
38532 maximum or minimum, respectively, for the scheduling policy specified by *policy*.

38533 The value of *policy* shall be one of the scheduling policy values defined in **<sched.h>**.

38534 **RETURN VALUE**

38535 If successful, the *sched\_get\_priority\_max()* and *sched\_get\_priority\_min()* functions shall return the  
38536 appropriate maximum or minimum values, respectively. If unsuccessful, they shall return a  
38537 value of  $-1$  and set *errno* to indicate the error.

38538 **ERRORS**

38539 The *sched\_get\_priority\_max()* and *sched\_get\_priority\_min()* functions shall fail if:

38540 [EINVAL] The value of the *policy* parameter does not represent a defined scheduling  
38541 policy.

38542 **EXAMPLES**

38543 None.

38544 **APPLICATION USAGE**

38545 None.

38546 **RATIONALE**

38547 None.

38548 **FUTURE DIRECTIONS**

38549 None.

38550 **SEE ALSO**

38551 *sched\_getparam()*, *sched\_setparam()*, *sched\_getscheduler()*, *sched\_rr\_get\_interval()*,  
38552 *sched\_setscheduler()*, the Base Definitions volume of IEEE Std 1003.1-200x, **<sched.h>**

38553 **CHANGE HISTORY**

38554 First released in Issue 5. Included for alignment with the POSIX Realtime Extension.

38555 **Issue 6**

38556 These functions are marked as part of the Process Scheduling option.

38557 The [ENOSYS] error condition has been removed as stubs need not be provided if an  
38558 implementation does not support the Process Scheduling option.

38559 The [ESRCH] error condition has been removed since these functions do not take a *pid*  
38560 argument.

38561 **NAME**38562 sched\_getparam — get scheduling parameters (**REALTIME**)38563 **SYNOPSIS**

38564 PS #include &lt;sched.h&gt;

38565 int sched\_getparam(pid\_t pid, struct sched\_param \*param);

38566

38567 **DESCRIPTION**38568 The *sched\_getparam()* function shall return the scheduling parameters of a process specified by  
38569 *pid* in the **sched\_param** structure pointed to by *param*.38570 If a process specified by *pid* exists, and if the calling process has permission, the scheduling  
38571 parameters for the process whose process ID is equal to *pid* shall be returned.38572 If *pid* is zero, the scheduling parameters for the calling process shall be returned. The behavior of  
38573 the *sched\_getparam()* function is unspecified if the value of *pid* is negative.38574 **RETURN VALUE**38575 Upon successful completion, the *sched\_getparam()* function shall return zero. If the call to  
38576 *sched\_getparam()* is unsuccessful, the function shall return a value of -1 and set *errno* to indicate  
38577 the error.38578 **ERRORS**38579 The *sched\_getparam()* function shall fail if:38580 [EPERM] The requesting process does not have permission to obtain the scheduling  
38581 parameters of the specified process.38582 [ESRCH] No process can be found corresponding to that specified by *pid*.38583 **EXAMPLES**

38584 None.

38585 **APPLICATION USAGE**

38586 None.

38587 **RATIONALE**

38588 None.

38589 **FUTURE DIRECTIONS**

38590 None.

38591 **SEE ALSO**38592 *sched\_getscheduler()*, *sched\_setparam()*, *sched\_setscheduler()*, the Base Definitions volume of  
38593 IEEE Std 1003.1-200x, <sched.h>38594 **CHANGE HISTORY**

38595 First released in Issue 5. Included for alignment with the POSIX Realtime Extension.

38596 **Issue 6**38597 The *sched\_getparam()* function is marked as part of the Process Scheduling option.38598 The [ENOSYS] error condition has been removed as stubs need not be provided if an  
38599 implementation does not support the Process Scheduling option.

38600 **NAME**

38601 sched\_getscheduler — get scheduling policy (**REALTIME**)

38602 **SYNOPSIS**

38603 PS #include <sched.h>

38604 int sched\_getscheduler(pid\_t pid);

38605

38606 **DESCRIPTION**

38607 The *sched\_getscheduler()* function shall return the scheduling policy of the process specified by  
 38608 *pid*. If the value of *pid* is negative, the behavior of the *sched\_getscheduler()* function is  
 38609 unspecified.

38610 The values that can be returned by *sched\_getscheduler()* are defined in the <**sched.h**> header.

38611 If a process specified by *pid* exists, and if the calling process has permission, the scheduling  
 38612 policy shall be returned for the process whose process ID is equal to *pid*.

38613 If *pid* is zero, the scheduling policy shall be returned for the calling process.

38614 **RETURN VALUE**

38615 Upon successful completion, the *sched\_getscheduler()* function shall return the scheduling policy  
 38616 of the specified process. If unsuccessful, the function shall return  $-1$  and set *errno* to indicate the  
 38617 error.

38618 **ERRORS**

38619 The *sched\_getscheduler()* function shall fail if:

38620 [EPERM] The requesting process does not have permission to determine the scheduling  
 38621 policy of the specified process.

38622 [ESRCH] No process can be found corresponding to that specified by *pid*.

38623 **EXAMPLES**

38624 None.

38625 **APPLICATION USAGE**

38626 None.

38627 **RATIONALE**

38628 None.

38629 **FUTURE DIRECTIONS**

38630 None.

38631 **SEE ALSO**

38632 *sched\_getparam()*, *sched\_setparam()*, *sched\_setscheduler()*, the Base Definitions volume of  
 38633 IEEE Std 1003.1-200x, <**sched.h**>

38634 **CHANGE HISTORY**

38635 First released in Issue 5. Included for alignment with the POSIX Realtime Extension.

38636 **Issue 6**

38637 The *sched\_getscheduler()* function is marked as part of the Process Scheduling option.

38638 The [ENOSYS] error condition has been removed as stubs need not be provided if an  
 38639 implementation does not support the Process Scheduling option.



38640 **NAME**

38641 sched\_rr\_get\_interval — get execution time limits (**REALTIME**)

38642 **SYNOPSIS**

38643 PS #include <sched.h>

38644 int sched\_rr\_get\_interval(pid\_t pid, struct timespec \*interval);

38645

38646 **DESCRIPTION**

38647 The *sched\_rr\_get\_interval()* function shall update the **timespec** structure referenced by the  
38648 *interval* argument to contain the current execution time limit (that is, time quantum) for the  
38649 process specified by *pid*. If *pid* is zero, the current execution time limit for the calling process  
38650 shall be returned.

38651 **RETURN VALUE**

38652 If successful, the *sched\_rr\_get\_interval()* function shall return zero. Otherwise, it shall return a  
38653 value of  $-1$  and set *errno* to indicate the error.

38654 **ERRORS**

38655 The *sched\_rr\_get\_interval()* function shall fail if:

38656 [ESRCH] No process can be found corresponding to that specified by *pid*.

38657 **EXAMPLES**

38658 None.

38659 **APPLICATION USAGE**

38660 None.

38661 **RATIONALE**

38662 None.

38663 **FUTURE DIRECTIONS**

38664 None.

38665 **SEE ALSO**

38666 *sched\_getparam()*, *sched\_get\_priority\_max()*, *sched\_getscheduler()*, *sched\_setparam()*,  
38667 *sched\_setscheduler()*, the Base Definitions volume of IEEE Std 1003.1-200x, <sched.h>

38668 **CHANGE HISTORY**

38669 First released in Issue 5. Included for alignment with the POSIX Realtime Extension.

38670 **Issue 6**

38671 The *sched\_rr\_get\_interval()* function is marked as part of the Process Scheduling option.

38672 The [ENOSYS] error condition has been removed as stubs need not be provided if an  
38673 implementation does not support the Process Scheduling option.

38674 **NAME**

38675 sched\_setparam — set scheduling parameters (**REALTIME**)

38676 **SYNOPSIS**

38677 PS `#include <sched.h>`

38678 `int sched_setparam(pid_t pid, const struct sched_param *param);`

38679

38680 **DESCRIPTION**

38681 The *sched\_setparam()* function shall set the scheduling parameters of the process specified by *pid*  
 38682 to the values specified by the **sched\_param** structure pointed to by *param*. The value of the  
 38683 *sched\_priority* member in the **sched\_param** structure shall be any integer within the inclusive  
 38684 priority range for the current scheduling policy of the process specified by *pid*. Higher  
 38685 numerical values for the priority represent higher priorities. If the value of *pid* is negative, the  
 38686 behavior of the *sched\_setparam()* function is unspecified.

38687 If a process specified by *pid* exists, and if the calling process has permission, the scheduling  
 38688 parameters shall be set for the process whose process ID is equal to *pid*.

38689 If *pid* is zero, the scheduling parameters shall be set for the calling process.

38690 The conditions under which one process has permission to change the scheduling parameters of  
 38691 another process are implementation-defined.

38692 Implementations may require the requesting process to have the appropriate privilege to set its  
 38693 own scheduling parameters or those of another process.

38694 The target process, whether it is running or not running, shall be moved to the tail of the thread  
 38695 list for its priority.

38696 If the priority of the process specified by the *pid* argument is set higher than that of the lowest  
 38697 priority running process and if the specified process is ready to run, the process specified by the  
 38698 *pid* argument shall preempt a lowest priority running process. Similarly, if the process calling  
 38699 *sched\_setparam()* sets its own priority lower than that of one or more other non-empty process  
 38700 lists, then the process that is the head of the highest priority list shall also preempt the calling  
 38701 process. Thus, in either case, the originating process might not receive notification of the  
 38702 completion of the requested priority change until the higher priority process has executed.

38703 ss If the scheduling policy of the target process is SCHED\_SPORADIC, the value specified by the  
 38704 *sched\_ss\_low\_priority* member of the *param* argument shall be any integer within the inclusive  
 38705 priority range for the sporadic server policy. The *sched\_ss\_repl\_period* and *sched\_ss\_init\_budget*  
 38706 members of the *param* argument shall represent the time parameters to be used by the sporadic  
 38707 server scheduling policy for the target process. The *sched\_ss\_max\_repl* member of the *param*  
 38708 argument shall represent the maximum number of replenishments that are allowed to be  
 38709 pending simultaneously for the process scheduled under this scheduling policy.

38710 The specified *sched\_ss\_repl\_period* shall be greater than or equal to the specified  
 38711 *sched\_ss\_init\_budget* for the function to succeed; if it is not, then the function shall fail.

38712 The value of *sched\_ss\_max\_repl* shall be within the inclusive range [1,{SS\_REPL\_MAX}] for the  
 38713 function to succeed; if not, the function shall fail.

38714 If the scheduling policy of the target process is either SCHED\_FIFO or SCHED\_RR, the  
 38715 *sched\_ss\_low\_priority*, *sched\_ss\_repl\_period*, and *sched\_ss\_init\_budget* members of the *param*  
 38716 argument shall have no effect on the scheduling behavior. If the scheduling policy of this process  
 38717 is not SCHED\_FIFO, SCHED\_RR, or SCHED\_SPORADIC, the effects of these members are  
 38718 implementation-defined; this case includes the SCHED\_OTHER policy.

38719 If the current scheduling policy for the process specified by *pid* is not SCHED\_FIFO,  
 38720 ss SCHED\_RR, or SCHED\_SPORADIC, the result is implementation-defined; this case includes the  
 38721 SCHED\_OTHER policy.

38722 The effect of this function on individual threads is dependent on the scheduling contention  
 38723 scope of the threads:

- 38724 • For threads with system scheduling contention scope, these functions shall have no effect on  
 38725 their scheduling.
- 38726 • For threads with process scheduling contention scope, the threads' scheduling parameters  
 38727 shall not be affected. However, the scheduling of these threads with respect to threads in  
 38728 other processes may be dependent on the scheduling parameters of their process, which are  
 38729 governed using these functions.

38730 If an implementation supports a two-level scheduling model in which library threads are  
 38731 multiplexed on top of several kernel-scheduled entities, then the underlying kernel-scheduled  
 38732 entities for the system contention scope threads shall not be affected by these functions.

38733 The underlying kernel-scheduled entities for the process contention scope threads shall have  
 38734 their scheduling parameters changed to the value specified in *param*. Kernel scheduled entities  
 38735 for use by process contention scope threads that are created after this call completes shall inherit  
 38736 their scheduling policy and associated scheduling parameters from the process.

38737 This function is not atomic with respect to other threads in the process. Threads may continue to  
 38738 execute while this function call is in the process of changing the scheduling policy for the  
 38739 underlying kernel-scheduled entities used by the process contention scope threads.

38740 **RETURN VALUE**

38741 If successful, the *sched\_setparam()* function shall return zero.

38742 If the call to *sched\_setparam()* is unsuccessful, the priority shall remain unchanged, and the  
 38743 function shall return a value of -1 and set *errno* to indicate the error.

38744 **ERRORS**

38745 The *sched\_setparam()* function shall fail if:

38746 38747	[EINVAL]	One or more of the requested scheduling parameters is outside the range defined for the scheduling policy of the specified <i>pid</i> .
38748 38749 38750	[EPERM]	The requesting process does not have permission to set the scheduling parameters for the specified process, or does not have the appropriate privilege to invoke <i>sched_setparam()</i> .
38751	[ESRCH]	No process can be found corresponding to that specified by <i>pid</i> .

38752 **EXAMPLES**

38753 None.

38754 **APPLICATION USAGE**

38755 None.

38756 **RATIONALE**

38757 None.

38758 **FUTURE DIRECTIONS**

38759 None.

38760 **SEE ALSO**

38761 *sched\_getparam()*, *sched\_getscheduler()*, *sched\_setscheduler()*, the Base Definitions volume of  
38762 IEEE Std 1003.1-200x, <**sched.h**>

38763 **CHANGE HISTORY**

38764 First released in Issue 5. Included for alignment with the POSIX Realtime Extension.

38765 **Issue 6**

38766 The *sched\_setparam()* function is marked as part of the Process Scheduling option.

38767 The [ENOSYS] error condition has been removed as stubs need not be provided if an  
38768 implementation does not support the Process Scheduling option.

38769 The following new requirements on POSIX implementations derive from alignment with the  
38770 Single UNIX Specification:

38771 • In the DESCRIPTION, the effect of this function on a thread's scheduling parameters is  
38772 added.

38773 • Sections describing two-level scheduling and atomicity of the function are added.

38774 The SCHED\_SPORADIC scheduling policy is added for alignment with IEEE Std 1003.1d-1999.

38775 IEEE PASC Interpretation 1003.1 #100 is applied.

## 38776 NAME

38777 sched\_setscheduler — set scheduling policy and parameters (**REALTIME**)

## 38778 SYNOPSIS

38779 PS 

```
#include <sched.h>
```

38780 

```
int sched_setscheduler(pid_t pid, int policy,
```

  
38781 

```
    const struct sched_param *param);
```

38782

## 38783 DESCRIPTION

38784 The *sched\_setscheduler()* function shall set the scheduling policy and scheduling parameters of  
 38785 the process specified by *pid* to *policy* and the parameters specified in the **sched\_param** structure  
 38786 pointed to by *param*, respectively. The value of the *sched\_priority* member in the **sched\_param**  
 38787 structure shall be any integer within the inclusive priority range for the scheduling policy  
 38788 specified by *policy*. If the value of *pid* is negative, the behavior of the *sched\_setscheduler()*  
 38789 function is unspecified.

38790 The possible values for the *policy* parameter are defined in the **<sched.h>** header.

38791 If a process specified by *pid* exists, and if the calling process has permission, the scheduling  
 38792 policy and scheduling parameters shall be set for the process whose process ID is equal to *pid*.

38793 If *pid* is zero, the scheduling policy and scheduling parameters shall be set for the calling  
 38794 process.

38795 The conditions under which one process has the appropriate privilege to change the scheduling  
 38796 parameters of another process are implementation-defined.

38797 Implementations may require that the requesting process have permission to set its own  
 38798 scheduling parameters or those of another process. Additionally, implementation-defined  
 38799 restrictions may apply as to the appropriate privileges required to set a process' own scheduling  
 38800 policy, or another process' scheduling policy, to a particular value.

38801 The *sched\_setscheduler()* function shall be considered successful if it succeeds in setting the  
 38802 scheduling policy and scheduling parameters of the process specified by *pid* to the values  
 38803 specified by *policy* and the structure pointed to by *param*, respectively.

38804 ss If the scheduling policy specified by *policy* is **SCHED\_SPORADIC**, the value specified by the  
 38805 *sched\_ss\_low\_priority* member of the *param* argument shall be any integer within the inclusive  
 38806 priority range for the sporadic server policy. The *sched\_ss\_repl\_period* and *sched\_ss\_init\_budget*  
 38807 members of the *param* argument shall represent the time parameters used by the sporadic server  
 38808 scheduling policy for the target process. The *sched\_ss\_max\_repl* member of the *param* argument  
 38809 shall represent the maximum number of replenishments that are allowed to be pending  
 38810 simultaneously for the process scheduled under this scheduling policy.

38811 The specified *sched\_ss\_repl\_period* shall be greater than or equal to the specified  
 38812 *sched\_ss\_init\_budget* for the function to succeed; if it is not, then the function shall fail.

38813 The value of *sched\_ss\_max\_repl* shall be within the inclusive range [1,{**SS\_REPL\_MAX**}] for the  
 38814 function to succeed; if not, the function shall fail.

38815 If the scheduling policy specified by *policy* is either **SCHED\_FIFO** or **SCHED\_RR**, the  
 38816 *sched\_ss\_low\_priority*, *sched\_ss\_repl\_period*, and *sched\_ss\_init\_budget* members of the *param*  
 38817 argument shall have no effect on the scheduling behavior.

38818 The effect of this function on individual threads is dependent on the scheduling contention  
 38819 scope of the threads:

38820       • For threads with system scheduling contention scope, these functions shall have no effect on  
38821 their scheduling.

38822       • For threads with process scheduling contention scope, the threads' scheduling policy and  
38823 associated parameters shall not be affected. However, the scheduling of these threads with  
38824 respect to threads in other processes may be dependent on the scheduling parameters of their  
38825 process, which are governed using these functions.

38826       If an implementation supports a two-level scheduling model in which library threads are  
38827 multiplexed on top of several kernel-scheduled entities, then the underlying kernel-scheduled  
38828 entities for the system contention scope threads shall not be affected by these functions.

38829       The underlying kernel-scheduled entities for the process contention scope threads shall have  
38830 their scheduling policy and associated scheduling parameters changed to the values specified in  
38831 *policy* and *param*, respectively. Kernel scheduled entities for use by process contention scope  
38832 threads that are created after this call completes shall inherit their scheduling policy and  
38833 associated scheduling parameters from the process.

38834       This function is not atomic with respect to other threads in the process. Threads may continue to  
38835 execute while this function call is in the process of changing the scheduling policy and  
38836 associated scheduling parameters for the underlying kernel-scheduled entities used by the  
38837 process contention scope threads.

#### 38838 RETURN VALUE

38839       Upon successful completion, the function shall return the former scheduling policy of the  
38840 specified process. If the *sched\_setscheduler()* function fails to complete successfully, the policy  
38841 and scheduling parameters shall remain unchanged, and the function shall return a value of  $-1$   
38842 and set *errno* to indicate the error.

#### 38843 ERRORS

38844       The *sched\_setscheduler()* function shall fail if:

38845       [EINVAL]       The value of the *policy* parameter is invalid, or one or more of the parameters  
38846 contained in *param* is outside the valid range for the specified scheduling  
38847 policy.

38848       [EPERM]       The requesting process does not have permission to set either or both of the  
38849 scheduling parameters or the scheduling policy of the specified process.

38850       [ESRCH]       No process can be found corresponding to that specified by *pid*.

#### 38851 EXAMPLES

38852       None.

#### 38853 APPLICATION USAGE

38854       None.

#### 38855 RATIONALE

38856       None.

#### 38857 FUTURE DIRECTIONS

38858       None.

#### 38859 SEE ALSO

38860       *sched\_getparam()*, *sched\_getscheduler()*, *sched\_setparam()*, the Base Definitions volume of  
38861 IEEE Std 1003.1-200x, <*sched.h*>

38862 **CHANGE HISTORY**

38863 First released in Issue 5. Included for alignment with the POSIX Realtime Extension.

38864 **Issue 6**

38865 The *sched\_setscheduler()* function is marked as part of the Process Scheduling option.

38866 The [ENOSYS] error condition has been removed as stubs need not be provided if an  
38867 implementation does not support the Process Scheduling option.

38868 The following new requirements on POSIX implementations derive from alignment with the  
38869 Single UNIX Specification:

38870 • In the DESCRIPTION, the effect of this function on a thread's scheduling parameters is  
38871 added.

38872 • Sections describing two-level scheduling and atomicity of the function are added.

38873 The SCHED\_SPORADIC scheduling policy is added for alignment with IEEE Std 1003.1d-1999.

38874 **NAME**

38875 sched\_yield — yield processor

38876 **SYNOPSIS**

38877 PS|THR #include &lt;sched.h&gt;

38878 int sched\_yield(void);

38879

38880 **DESCRIPTION**38881 The *sched\_yield()* function shall force the running thread to relinquish the processor until it again  
38882 becomes the head of its thread list. It takes no arguments.38883 **RETURN VALUE**38884 The *sched\_yield()* function shall return 0 if it completes successfully; otherwise, it shall return a  
38885 value of -1 and set *errno* to indicate the error.38886 **ERRORS**

38887 No errors are defined.

38888 **EXAMPLES**

38889 None.

38890 **APPLICATION USAGE**

38891 None.

38892 **RATIONALE**

38893 None.

38894 **FUTURE DIRECTIONS**

38895 None.

38896 **SEE ALSO**

38897 The Base Definitions volume of IEEE Std 1003.1-200x, &lt;sched.h&gt;

38898 **CHANGE HISTORY**38899 First released in Issue 5. Included for alignment with the POSIX Realtime Extension and the  
38900 POSIX Threads Extension.38901 **Issue 6**38902 The *sched\_yield()* function is now marked as part of the Process Scheduling and Threads options.



38903 **NAME**

38904           seed48 — seed uniformly distributed pseudo-random non-negative long integer generator

38905 **SYNOPSIS**

38906 xSI       #include <stdlib.h>

38907           unsigned short \*seed48(unsigned short seed16v[3]);

38908

38909 **DESCRIPTION**

38910           Refer to *drand48()*.

38911 **NAME**

38912 seekdir — set position of directory stream

38913 **SYNOPSIS**

38914 XSI #include &lt;dirent.h&gt;

38915 void seekdir(DIR \*dirp, long loc);

38916

38917 **DESCRIPTION**

38918 The *seekdir()* function shall set the position of the next *readdir()* operation on the directory  
38919 stream specified by *dirp* to the position specified by *loc*. The value of *loc* should have been  
38920 returned from an earlier call to *telldir()*. The new position reverts to the one associated with the  
38921 directory stream when *telldir()* was performed.

38922 If the value of *loc* was not obtained from an earlier call to *telldir()*, or if a call to *rewinddir()*  
38923 occurred between the call to *telldir()* and the call to *seekdir()*, the results of subsequent calls to  
38924 *readdir()* are unspecified.

38925 **RETURN VALUE**38926 The *seekdir()* function shall not return a value.38927 **ERRORS**

38928 No errors are defined.

38929 **EXAMPLES**

38930 None.

38931 **APPLICATION USAGE**

38932 None.

38933 **RATIONALE**

38934 The original standard developers perceived that there were restrictions on the use of the  
38935 *seekdir()* and *telldir()* functions related to implementation details, and for that reason these  
38936 functions need not be supported on all POSIX-conforming systems. They are required on  
38937 implementations supporting the XSI extension.

38938 One of the perceived problems of implementation is that returning to a given point in a directory  
38939 is quite difficult to describe formally, in spite of its intuitive appeal, when systems that use B-  
38940 trees, hashing functions, or other similar mechanisms to order their directories are considered.  
38941 The definition of *seekdir()* and *telldir()* does not specify whether, when using these interfaces, a  
38942 given directory entry will be seen at all, or more than once.

38943 On systems not supporting these functions, their capability can sometimes be accomplished by  
38944 saving a filename found by *readdir()* and later using *rewinddir()* and a loop on *readdir()* to  
38945 relocate the position from which the filename was saved.

38946 **FUTURE DIRECTIONS**

38947 None.

38948 **SEE ALSO**

38949 *opendir()*, *readdir()*, *telldir()*, the Base Definitions volume of IEEE Std 1003.1-200x, <dirent.h>,  
38950 <stdio.h>, <sys/types.h>

38951 **CHANGE HISTORY**

38952 First released in Issue 2.

38953 **Issue 6**

38954 In the SYNOPSIS, the inclusion of `<sys/types.h>` is no longer required.

38955 **NAME**

38956           select — synchronous I/O multiplexing

38957 **SYNOPSIS**

38958           #include <sys/time.h>

```
38959           int select(int nfds, fd_set *restrict readfds,  
38960                      fd_set *restrict writefds, fd_set *restrict errorfds,  
38961                      struct timeval *restrict timeout);
```

38962

38963 **DESCRIPTION**

38964           Refer to *pselect()*.

38965 **NAME**38966 sem\_close — close a named semaphore (**REALTIME**)38967 **SYNOPSIS**

38968 SEM #include &lt;semaphore.h&gt;

38969 int sem\_close(sem\_t \*sem);

38970

38971 **DESCRIPTION**

38972 The `sem_close()` function shall indicate that the calling process is finished using the named  
 38973 semaphore indicated by `sem`. The effects of calling `sem_close()` for an unnamed semaphore (one  
 38974 created by `sem_init()`) are undefined. The `sem_close()` function shall deallocate (that is, make  
 38975 available for reuse by a subsequent `sem_open()` by this process) any system resources allocated  
 38976 by the system for use by this process for this semaphore. The effect of subsequent use of the  
 38977 semaphore indicated by `sem` by this process is undefined. If the semaphore has not been  
 38978 removed with a successful call to `sem_unlink()`, then `sem_close()` has no effect on the state of the  
 38979 semaphore. If the `sem_unlink()` function has been successfully invoked for `name` after the most  
 38980 recent call to `sem_open()` with `O_CREAT` for this semaphore, then when all processes that have  
 38981 opened the semaphore close it, the semaphore is no longer accessible.

38982 **RETURN VALUE**

38983 Upon successful completion, a value of zero shall be returned. Otherwise, a value of `-1` shall be  
 38984 returned and `errno` set to indicate the error.

38985 **ERRORS**38986 The `sem_close()` function shall fail if:38987 [EINVAL] The `sem` argument is not a valid semaphore descriptor.38988 **EXAMPLES**

38989 None.

38990 **APPLICATION USAGE**

38991 The `sem_close()` function is part of the Semaphores option and need not be available on all  
 38992 implementations.

38993 **RATIONALE**

38994 None.

38995 **FUTURE DIRECTIONS**

38996 None.

38997 **SEE ALSO**

38998 `semctl()`, `semget()`, `semop()`, `sem_init()`, `sem_open()`, `sem_unlink()`, the Base Definitions volume of  
 38999 IEEE Std 1003.1-200x, <**semaphore.h**>

39000 **CHANGE HISTORY**

39001 First released in Issue 5. Included for alignment with the POSIX Realtime Extension.

39002 **Issue 6**39003 The `sem_close()` function is marked as part of the Semaphores option.

39004 The [ENOSYS] error condition has been removed as stubs need not be provided if an  
 39005 implementation does not support the Semaphores option.

39006 **NAME**

39007 sem\_destroy — destroy an unnamed semaphore (**REALTIME**)

39008 **SYNOPSIS**

39009 SEM #include <semaphore.h>

39010 int sem\_destroy(sem\_t \*sem);

39011

39012 **DESCRIPTION**

39013 The *sem\_destroy()* function shall destroy the unnamed semaphore indicated by *sem*. Only a  
39014 semaphore that was created using *sem\_init()* may be destroyed using *sem\_destroy()*; the effect of  
39015 calling *sem\_destroy()* with a named semaphore is undefined. The effect of subsequent use of the  
39016 semaphore *sem* is undefined until *sem* is reinitialized by another call to *sem\_init()*.

39017 It is safe to destroy an initialized semaphore upon which no threads are currently blocked. The  
39018 effect of destroying a semaphore upon which other threads are currently blocked is undefined.

39019 **RETURN VALUE**

39020 Upon successful completion, a value of zero shall be returned. Otherwise, a value of -1 shall be  
39021 returned and *errno* set to indicate the error.

39022 **ERRORS**

39023 The *sem\_destroy()* function shall fail if:

39024 [EINVAL] The *sem* argument is not a valid semaphore.

39025 The *sem\_destroy()* function may fail if:

39026 [EBUSY] There are currently processes blocked on the semaphore.

39027 **EXAMPLES**

39028 None.

39029 **APPLICATION USAGE**

39030 The *sem\_destroy()* function is part of the Semaphores option and need not be available on all  
39031 implementations.

39032 **RATIONALE**

39033 None.

39034 **FUTURE DIRECTIONS**

39035 None.

39036 **SEE ALSO**

39037 *semctl()*, *semget()*, *semop()*, *sem\_init()*, *sem\_open()*, the Base Definitions volume of  
39038 IEEE Std 1003.1-200x, <semaphore.h>

39039 **CHANGE HISTORY**

39040 First released in Issue 5. Included for alignment with the POSIX Realtime Extension.

39041 **Issue 6**

39042 The *sem\_destroy()* function is marked as part of the Semaphores option.

39043 The [ENOSYS] error condition has been removed as stubs need not be provided if an  
39044 implementation does not support the Semaphores option.

39045 **NAME**

39046 `sem_getvalue` — get the value of a semaphore (**REALTIME**)

39047 **SYNOPSIS**

```
39048 SEM #include <semaphore.h>
```

```
39049 int sem_getvalue(sem_t *restrict sem, int *restrict sval);
```

39050

39051 **DESCRIPTION**

39052 The `sem_getvalue()` function shall update the location referenced by the `sval` argument to have  
39053 the value of the semaphore referenced by `sem` without affecting the state of the semaphore. The  
39054 updated value represents an actual semaphore value that occurred at some unspecified time  
39055 during the call, but it need not be the actual value of the semaphore when it is returned to the  
39056 calling process.

39057 If `sem` is locked, then the value returned by `sem_getvalue()` is either zero or a negative number  
39058 whose absolute value represents the number of processes waiting for the semaphore at some  
39059 unspecified time during the call.

39060 **RETURN VALUE**

39061 Upon successful completion, the `sem_getvalue()` function shall return a value of zero. Otherwise,  
39062 it shall return a value of `-1` and set `errno` to indicate the error.

39063 **ERRORS**

39064 The `sem_getvalue()` function shall fail if:

39065 [EINVAL] The `sem` argument does not refer to a valid semaphore.

39066 **EXAMPLES**

39067 None.

39068 **APPLICATION USAGE**

39069 The `sem_getvalue()` function is part of the Semaphores option and need not be available on all  
39070 implementations.

39071 **RATIONALE**

39072 None.

39073 **FUTURE DIRECTIONS**

39074 None.

39075 **SEE ALSO**

39076 `semctl()`, `semget()`, `semop()`, `sem_post()`, `sem_timedwait()`, `sem_trywait()`, `sem_wait()`, the Base  
39077 Definitions volume of IEEE Std 1003.1-200x, <**semaphore.h**>

39078 **CHANGE HISTORY**

39079 First released in Issue 5. Included for alignment with the POSIX Realtime Extension.

39080 **Issue 6**

39081 The `sem_getvalue()` function is marked as part of the Semaphores option.

39082 The [ENOSYS] error condition has been removed as stubs need not be provided if an  
39083 implementation does not support the Semaphores option.

39084 The `sem_timedwait()` function is added to the SEE ALSO section for alignment with  
39085 IEEE Std 1003.1d-1999.

39086 The **restrict** keyword is added to the `sem_getvalue()` prototype for alignment with the  
39087 ISO/IEC 9899:1999 standard.

39088 **NAME**39089 sem\_init — initialize an unnamed semaphore (**REALTIME**)39090 **SYNOPSIS**

39091 SEM #include &lt;semaphore.h&gt;

39092 int sem\_init(sem\_t \*sem, int pshared, unsigned value);

39093

39094 **DESCRIPTION**

39095 The *sem\_init()* function shall initialize the unnamed semaphore referred to by *sem*. The value of  
 39096 the initialized semaphore shall be *value*. Following a successful call to *sem\_init()*, the semaphore  
 39097 may be used in subsequent calls to *sem\_wait()*, *sem\_trywait()*, *sem\_post()*, and *sem\_destroy()*.  
 39098 This semaphore shall remain usable until the semaphore is destroyed.

39099 If the *pshared* argument has a non-zero value, then the semaphore is shared between processes;  
 39100 in this case, any process that can access the semaphore *sem* can use *sem* for performing  
 39101 *sem\_wait()*, *sem\_trywait()*, *sem\_post()*, and *sem\_destroy()* operations.

39102 Only *sem* itself may be used for performing synchronization. The result of referring to copies of  
 39103 *sem* in calls to *sem\_wait()*, *sem\_trywait()*, *sem\_post()*, and *sem\_destroy()*, is undefined.

39104 If the *pshared* argument is zero, then the semaphore is shared between threads of the process; any  
 39105 thread in this process can use *sem* for performing *sem\_wait()*, *sem\_trywait()*, *sem\_post()*, and  
 39106 *sem\_destroy()* operations. The use of the semaphore by threads other than those created in the  
 39107 same process is undefined.

39108 Attempting to initialize an already initialized semaphore results in undefined behavior.

39109 **RETURN VALUE**

39110 Upon successful completion, the *sem\_init()* function shall initialize the semaphore in *sem*.  
 39111 Otherwise, it shall return  $-1$  and set *errno* to indicate the error.

39112 **ERRORS**39113 The *sem\_init()* function shall fail if:39114 [EINVAL] The *value* argument exceeds {SEM\_VALUE\_MAX}.

39115 [ENOSPC] A resource required to initialize the semaphore has been exhausted, or the  
 39116 limit on semaphores ({SEM\_NSEMS\_MAX}) has been reached.

39117 [EPERM] The process lacks the appropriate privileges to initialize the semaphore.

39118 **EXAMPLES**

39119 None.

39120 **APPLICATION USAGE**

39121 The *sem\_init()* function is part of the Semaphores option and need not be available on all  
 39122 implementations.

39123 **RATIONALE**

39124 Although this volume of IEEE Std 1003.1-200x fails to specify a successful return value, it is  
 39125 likely that a later version may require the implementation to return a value of zero if the call to  
 39126 *sem\_init()* is successful.

39127 **FUTURE DIRECTIONS**

39128 None.



39129 **SEE ALSO**

39130 *sem\_destroy()*, *sem\_post()*, *sem\_timedwait()*, *sem\_trywait()*, *sem\_wait()*, the Base Definitions  
39131 volume of IEEE Std 1003.1-200x, <semaphore.h>

39132 **CHANGE HISTORY**

39133 First released in Issue 5. Included for alignment with the POSIX Realtime Extension.

39134 **Issue 6**

39135 The *sem\_init()* function is marked as part of the Semaphores option.

39136 The [ENOSYS] error condition has been removed as stubs need not be provided if an  
39137 implementation does not support the Semaphores option.

39138 The *sem\_timedwait()* function is added to the SEE ALSO section for alignment with  
39139 IEEE Std 1003.1d-1999.

## 39140 NAME

39141 sem\_open — initialize and open a named semaphore (**REALTIME**)

## 39142 SYNOPSIS

39143 SEM #include &lt;semaphore.h&gt;

39144 sem\_t \*sem\_open(const char \*name, int oflag, ...);

39145

## 39146 DESCRIPTION

39147 The *sem\_open()* function shall establish a connection between a named semaphore and a process.  
 39148 Following a call to *sem\_open()* with semaphore name *name*, the process may reference the  
 39149 semaphore associated with *name* using the address returned from the call. This semaphore may  
 39150 be used in subsequent calls to *sem\_wait()*, *sem\_trywait()*, *sem\_post()*, and *sem\_close()*. The  
 39151 semaphore remains usable by this process until the semaphore is closed by a successful call to  
 39152 *sem\_close()*, *\_exit()*, or one of the *exec* functions.

39153 The *oflag* argument controls whether the semaphore is created or merely accessed by the call to  
 39154 *sem\_open()*. The following flag bits may be set in *oflag*:

39155 **O\_CREAT** This flag is used to create a semaphore if it does not already exist. If **O\_CREAT**  
 39156 is set and the semaphore already exists, then **O\_CREAT** has no effect, except as noted  
 39157 under **O\_EXCL**. Otherwise, *sem\_open()* creates a named semaphore. The **O\_CREAT**  
 39158 flag requires a third and a fourth argument: *mode*, which is of type **mode\_t**, and  
 39159 *value*, which is of type **unsigned**. The semaphore is created with an initial value of  
 39160 *value*. Valid initial values for semaphores are less than or equal to  
 39161 {SEM\_VALUE\_MAX}.

39162 The user ID of the semaphore is set to the effective user ID of the process; the  
 39163 group ID of the semaphore is set to a system default group ID or to the effective  
 39164 group ID of the process. The permission bits of the semaphore are set to the value  
 39165 of the *mode* argument except those set in the file mode creation mask of the  
 39166 process. When bits in *mode* other than the file permission bits are specified, the  
 39167 effect is unspecified.

39168 After the semaphore named *name* has been created by *sem\_open()* with the  
 39169 **O\_CREAT** flag, other processes can connect to the semaphore by calling  
 39170 *sem\_open()* with the same value of *name*.

39171 **O\_EXCL** If **O\_EXCL** and **O\_CREAT** are set, *sem\_open()* fails if the semaphore *name* exists.  
 39172 The check for the existence of the semaphore and the creation of the semaphore if  
 39173 it does not exist are atomic with respect to other processes executing *sem\_open()*  
 39174 with **O\_EXCL** and **O\_CREAT** set. If **O\_EXCL** is set and **O\_CREAT** is not set, the  
 39175 effect is undefined.

39176 If flags other than **O\_CREAT** and **O\_EXCL** are specified in the *oflag* parameter, the  
 39177 effect is unspecified.

39178 The *name* argument points to a string naming a semaphore object. It is unspecified whether the  
 39179 name appears in the file system and is visible to functions that take pathnames as arguments. |  
 39180 The *name* argument conforms to the construction rules for a pathname. If *name* begins with the |  
 39181 slash character, then processes calling *sem\_open()* with the same value of *name* shall refer to the |  
 39182 same semaphore object, as long as that name has not been removed. If *name* does not begin with |  
 39183 the slash character, the effect is implementation-defined. The interpretation of slash characters |  
 39184 other than the leading slash character in *name* is implementation-defined.

39185 If a process makes multiple successful calls to *sem\_open()* with the same value for *name*, the |  
 39186 same semaphore address shall be returned for each such successful call, provided that there |

- 39187 have been no calls to *sem\_unlink()* for this semaphore.
- 39188 References to copies of the semaphore produce undefined results.
- 39189 **RETURN VALUE**
- 39190 Upon successful completion, the *sem\_open()* function shall return the address of the semaphore.
- 39191 Otherwise, it shall return a value of SEM\_FAILED and set *errno* to indicate the error. The symbol
- 39192 SEM\_FAILED is defined in the <**semaphore.h**> header. No successful return from *sem\_open()*
- 39193 shall return the value SEM\_FAILED.
- 39194 **ERRORS**
- 39195 If any of the following conditions occur, the *sem\_open()* function shall return SEM\_FAILED and
- 39196 set *errno* to the corresponding value:
- 39197 [EACCES] The named semaphore exists and the permissions specified by *oflag* are
- 39198 denied, or the named semaphore does not exist and permission to create the
- 39199 named semaphore is denied.
- 39200 [EEXIST] O\_CREAT and O\_EXCL are set and the named semaphore already exists.
- 39201 [EINTR] The *sem\_open()* operation was interrupted by a signal.
- 39202 [EINVAL] The *sem\_open()* operation is not supported for the given name, or O\_CREAT
- 39203 was specified in *oflag* and *value* was greater than {SEM\_VALUE\_MAX}.
- 39204 [EMFILE] Too many semaphore descriptors or file descriptors are currently in use by
- 39205 this process.
- 39206 [ENAMETOOLONG]
- 39207 The length of the *name* argument exceeds {PATH\_MAX} or a pathname
- 39208 component is longer than {NAME\_MAX}.
- 39209 [ENFILE] Too many semaphores are currently open in the system.
- 39210 [ENOENT] O\_CREAT is not set and the named semaphore does not exist.
- 39211 [ENOSPC] There is insufficient space for the creation of the new named semaphore.
- 39212 **EXAMPLES**
- 39213 None.
- 39214 **APPLICATION USAGE**
- 39215 The *sem\_open()* function is part of the Semaphores option and need not be available on all
- 39216 implementations.
- 39217 **RATIONALE**
- 39218 An earlier version of this volume of IEEE Std 1003.1-200x required an error return value of -1
- 39219 with the type **sem\_t \*** for the *sem\_open()* function, which is not guaranteed to be portable across
- 39220 implementations. The revised text provides the symbolic error code SEM\_FAILED to eliminate
- 39221 the type conflict.
- 39222 **FUTURE DIRECTIONS**
- 39223 None.
- 39224 **SEE ALSO**
- 39225 *semctl()*, *semget()*, *semop()*, *sem\_close()*, *sem\_post()*, *sem\_timedwait()*, *sem\_trywait()*, *sem\_unlink()*,
- 39226 *sem\_wait()*, the Base Definitions volume of IEEE Std 1003.1-200x, <**semaphore.h**>

39227 **CHANGE HISTORY**

39228 First released in Issue 5. Included for alignment with the POSIX Realtime Extension.

39229 **Issue 6**

39230 The *sem\_open()* function is marked as part of the Semaphores option.

39231 The [ENOSYS] error condition has been removed as stubs need not be provided if an  
39232 implementation does not support the Semaphores option.

39233 The *sem\_timedwait()* function is added to the SEE ALSO section for alignment with  
39234 IEEE Std 1003.1d-1999.

39235 **NAME**39236 sem\_post — unlock a semaphore (**REALTIME**)39237 **SYNOPSIS**

39238 SEM #include &lt;semaphore.h&gt;

39239 int sem\_post(sem\_t \*sem);

39240

39241 **DESCRIPTION**39242 The *sem\_post()* function shall unlock the semaphore referenced by *sem* by performing a  
39243 semaphore unlock operation on that semaphore.39244 If the semaphore value resulting from this operation is positive, then no threads were blocked  
39245 waiting for the semaphore to become unlocked; the semaphore value is simply incremented.39246 If the value of the semaphore resulting from this operation is zero, then one of the threads  
39247 blocked waiting for the semaphore shall be allowed to return successfully from its call to  
39248 *sem\_wait()*. If the Process Scheduling option is supported, the thread to be unblocked shall be  
39249 chosen in a manner appropriate to the scheduling policies and parameters in effect for the  
39250 blocked threads. In the case of the schedulers SCHED\_FIFO and SCHED\_RR, the highest  
39251 priority waiting thread shall be unblocked, and if there is more than one highest priority thread  
39252 blocked waiting for the semaphore, then the highest priority thread that has been waiting the  
39253 longest shall be unblocked. If the Process Scheduling option is not defined, the choice of a thread  
39254 to unblock is unspecified.39255 SS If the Process Sporadic Server option is supported, and the scheduling policy is  
39256 SCHED\_SPORADIC, the semantics are as per SCHED\_FIFO above.39257 The *sem\_post()* function shall be reentrant with respect to signals and may be invoked from a  
39258 signal-catching function.39259 **RETURN VALUE**39260 If successful, the *sem\_post()* function shall return zero; otherwise, the function shall return  $-1$   
39261 and set *errno* to indicate the error.39262 **ERRORS**39263 The *sem\_post()* function shall fail if:39264 [EINVAL] The *sem* argument does not refer to a valid semaphore.39265 **EXAMPLES**

39266 None.

39267 **APPLICATION USAGE**39268 The *sem\_post()* function is part of the Semaphores option and need not be available on all  
39269 implementations.39270 **RATIONALE**

39271 None.

39272 **FUTURE DIRECTIONS**

39273 None.

39274 **SEE ALSO**39275 *semctl()*, *semget()*, *semop()*, *sem\_timedwait()*, *sem\_trywait()*, *sem\_wait()*, the Base Definitions  
39276 volume of IEEE Std 1003.1-200x, <semaphore.h>

39277 **CHANGE HISTORY**

39278 First released in Issue 5. Included for alignment with the POSIX Realtime Extension.

39279 **Issue 6**

39280 The *sem\_post()* function is marked as part of the Semaphores option.

39281 The [ENOSYS] error condition has been removed as stubs need not be provided if an  
39282 implementation does not support the Semaphores option.

39283 The *sem\_timedwait()* function is added to the SEE ALSO section for alignment with  
39284 IEEE Std 1003.1d-1999.

39285 SCHED\_SPORADIC is added to the list of scheduling policies for which the thread that is to be  
39286 unblocked is specified for alignment with IEEE Std 1003.1d-1999.

39287 **NAME**39288 sem\_timedwait — lock a semaphore (**ADVANCED REALTIME**)39289 **SYNOPSIS**

39290 SEM TMO #include &lt;semaphore.h&gt;

39291 #include &lt;time.h&gt;

```
39292 int sem_timedwait(sem_t *restrict sem,
39293                  const struct timespec *restrict abs_timeout);
39294
```

39295 **DESCRIPTION**

39296 The *sem\_timedwait()* function shall lock the semaphore referenced by *sem* as in the *sem\_wait()*  
 39297 function. However, if the semaphore cannot be locked without waiting for another process or  
 39298 thread to unlock the semaphore by performing a *sem\_post()* function, this wait shall be  
 39299 terminated when the specified timeout expires.

39300 The timeout shall expire when the absolute time specified by *abs\_timeout* passes, as measured by  
 39301 the clock on which timeouts are based (that is, when the value of that clock equals or exceeds  
 39302 *abs\_timeout*), or if the absolute time specified by *abs\_timeout* has already been passed at the time  
 39303 of the call.

39304 TMR If the Timers option is supported, the timeout shall be based on the `CLOCK_REALTIME` clock. If  
 39305 the Timers option is not supported, the timeout shall be based on the system clock as returned  
 39306 by the *time()* function. The resolution of the timeout shall be the resolution of the clock on which  
 39307 it is based. The **timespec** data type is defined as a structure in the `<time.h>` header.

39308 Under no circumstance shall the function fail with a timeout if the semaphore can be locked  
 39309 immediately. The validity of the *abs\_timeout* need not be checked if the semaphore can be locked  
 39310 immediately.

39311 **RETURN VALUE**

39312 The *sem\_timedwait()* function shall return zero if the calling process successfully performed the  
 39313 semaphore lock operation on the semaphore designated by *sem*. If the call was unsuccessful, the  
 39314 state of the semaphore shall be unchanged, and the function shall return a value of `-1` and set  
 39315 *errno* to indicate the error.

39316 **ERRORS**39317 The *sem\_timedwait()* function shall fail if:39318 [EINVAL] The *sem* argument does not refer to a valid semaphore.

39319 [EINVAL] The process or thread would have blocked, and the *abs\_timeout* parameter  
 39320 specified a nanoseconds field value less than zero or greater than or equal to  
 39321 1 000 million.

39322 [ETIMEDOUT] The semaphore could not be locked before the specified timeout expired.

39323 The *sem\_timedwait()* function may fail if:

39324 [EDEADLK] A deadlock condition was detected.

39325 [EINTR] A signal interrupted this function.

**39326 EXAMPLES**

39327 None.

**39328 APPLICATION USAGE**

39329 Applications using these functions may be subject to priority inversion, as discussed in the Base  
39330 Definitions volume of IEEE Std 1003.1-200x, Section 3.285, Priority Inversion.

39331 The *sem\_timedwait()* function is part of the Semaphores and Timeouts options and need not be  
39332 provided on all implementations.

**39333 RATIONALE**

39334 None.

**39335 FUTURE DIRECTIONS**

39336 None.

**39337 SEE ALSO**

39338 *sem\_post()*, *sem\_trywait()*, *sem\_wait()*, *semctl()*, *semget()*, *semop()*, *time()*, the Base Definitions  
39339 volume of IEEE Std 1003.1-200x, <**semaphore.h**>, <**time.h**>

**39340 CHANGE HISTORY**

39341 First released in Issue 6. Derived from IEEE Std 1003.1d-1999.



39342 **NAME**39343 sem\_trywait, sem\_wait — lock a semaphore (**REALTIME**)39344 **SYNOPSIS**

39345 SEM #include &lt;semaphore.h&gt;

39346 int sem\_trywait(sem\_t \*sem);

39347 int sem\_wait(sem\_t \*sem);

39348

39349 **DESCRIPTION**

39350 The *sem\_trywait()* function shall lock the semaphore referenced by *sem* only if the semaphore is  
 39351 currently not locked; that is, if the semaphore value is currently positive. Otherwise, shall does  
 39352 not lock the semaphore.

39353 The *sem\_wait()* function shall lock the semaphore referenced by *sem* by performing a semaphore  
 39354 lock operation on that semaphore. If the semaphore value is currently zero, then the calling  
 39355 thread shall not return from the call to *sem\_wait()* until it either locks the semaphore or the call is  
 39356 interrupted by a signal.

39357 Upon successful return, the state of the semaphore shall be locked and shall remain locked until  
 39358 the *sem\_post()* function is executed and returns successfully.

39359 The *sem\_wait()* function is interruptible by the delivery of a signal.

39360 **RETURN VALUE**

39361 The *sem\_trywait()* and *sem\_wait()* functions shall return zero if the calling process successfully  
 39362 performed the semaphore lock operation on the semaphore designated by *sem*. If the call was  
 39363 unsuccessful, the state of the semaphore shall be unchanged, and the function shall return a  
 39364 value of  $-1$  and set *errno* to indicate the error.

39365 **ERRORS**

39366 The *sem\_trywait()* and *sem\_wait()* functions shall fail if:

39367 [EAGAIN] The semaphore was already locked, so it cannot be immediately locked by the  
 39368 *sem\_trywait()* operation (*sem\_trywait()* only).

39369 [EINVAL] The *sem* argument does not refer to a valid semaphore.

39370 The *sem\_trywait()* and *sem\_wait()* functions may fail if:

39371 [EDEADLK] A deadlock condition was detected.

39372 [EINTR] A signal interrupted this function.

39373 **EXAMPLES**

39374 None.

39375 **APPLICATION USAGE**

39376 Applications using these functions may be subject to priority inversion, as discussed in the Base  
 39377 Definitions volume of IEEE Std 1003.1-200x, Section 3.285, Priority Inversion.

39378 The *sem\_trywait()* and *sem\_wait()* functions are part of the Semaphores option and need not be  
 39379 provided on all implementations.

39380 **RATIONALE**

39381 None.

39382 **FUTURE DIRECTIONS**

39383 None.

39384 **SEE ALSO**

39385 *semctl()*, *semget()*, *semop()*, *sem\_post()*, *sem\_timedwait()*, the Base Definitions volume of  
39386 IEEE Std 1003.1-200x, <**semaphore.h**>

39387 **CHANGE HISTORY**

39388 First released in Issue 5. Included for alignment with the POSIX Realtime Extension.

39389 **Issue 6**39390 The *sem\_trywait()* and *sem\_wait()* functions are marked as part of the Semaphores option.

39391 The [ENOSYS] error condition has been removed as stubs need not be provided if an  
39392 implementation does not support the Semaphores option.

39393 The *sem\_timedwait()* function is added to the SEE ALSO section for alignment with  
39394 IEEE Std 1003.1d-1999.

39395 **NAME**39396 sem\_unlink — remove a named semaphore (**REALTIME**)39397 **SYNOPSIS**

39398 SEM #include &lt;semaphore.h&gt;

39399 int sem\_unlink(const char \*name);

39400

39401 **DESCRIPTION**

39402 The *sem\_unlink()* function shall remove the semaphore named by the string *name*. If the  
 39403 semaphore named by *name* is currently referenced by other processes, then *sem\_unlink()* shall  
 39404 have no effect on the state of the semaphore. If one or more processes have the semaphore open  
 39405 when *sem\_unlink()* is called, destruction of the semaphore is postponed until all references to the  
 39406 semaphore have been destroyed by calls to *sem\_close()*, *\_exit()*, or *exec*. Calls to *sem\_open()* to  
 39407 recreate or reconnect to the semaphore refer to a new semaphore after *sem\_unlink()* is called. The  
 39408 *sem\_unlink()* call shall not block until all references have been destroyed; it shall return  
 39409 immediately.

39410 **RETURN VALUE**

39411 Upon successful completion, the *sem\_unlink()* function shall return a value of 0. Otherwise, the  
 39412 semaphore shall not be changed and the function shall return a value of -1 and set *errno* to  
 39413 indicate the error.

39414 **ERRORS**39415 The *sem\_unlink()* function shall fail if:

39416 [EACCES] Permission is denied to unlink the named semaphore.

39417 [ENAMETOOLONG]

39418 The length of the *name* argument exceeds {PATH\_MAX} or a pathname  
 39419 component is longer than {NAME\_MAX}.

39420 [ENOENT] The named semaphore does not exist.

39421 **EXAMPLES**

39422 None.

39423 **APPLICATION USAGE**

39424 The *sem\_unlink()* function is part of the Semaphores option and need not be available on all  
 39425 implementations.

39426 **RATIONALE**

39427 None.

39428 **FUTURE DIRECTIONS**

39429 None.

39430 **SEE ALSO**

39431 *semctl()*, *semget()*, *semop()*, *sem\_close()*, *sem\_open()*, the Base Definitions volume of  
 39432 IEEE Std 1003.1-200x, <semaphore.h>

39433 **CHANGE HISTORY**

39434 First released in Issue 5. Included for alignment with the POSIX Realtime Extension.

39435 **Issue 6**39436 The *sem\_unlink()* function is marked as part of the Semaphores option.

39437  
39438

The [ENOSYS] error condition has been removed as stubs need not be provided if an implementation does not support the Semaphores option.

39439 **NAME**39440 sem\_wait — lock a semaphore (**REALTIME**)39441 **SYNOPSIS**

39442 SEM #include &lt;semaphore.h&gt;

39443 int sem\_wait(sem\_t \*sem);

39444

39445 **DESCRIPTION**39446 Refer to *sem\_trywait()*.

## 39447 NAME

39448 semctl — XSI semaphore control operations

## 39449 SYNOPSIS

39450 XSI 

```
#include <sys/sem.h>
```

39451 

```
int semctl(int semid, int semnum, int cmd, ...);
```

39452

## 39453 DESCRIPTION

39454 The *semctl()* function operates on XSI semaphores (see the Base Definitions volume of |  
 39455 IEEE Std 1003.1-200x, Section 4.15, Semaphore). It is unspecified whether this function |  
 39456 interoperates with the realtime interprocess communication facilities defined in Section 2.8 (on  
 39457 page 491).

39458 The *semctl()* function provides a variety of semaphore control operations as specified by *cmd*.  
 39459 The fourth argument is optional and depends upon the operation requested. If required, it is of  
 39460 type **union semun**, which the application shall explicitly declare:

```
39461 union semun {
39462     int val;
39463     struct semid_ds *buf;
39464     unsigned short *array;
39465 } arg;
```

39466 The following semaphore control operations as specified by *cmd* are executed with respect to the  
 39467 semaphore specified by *semid* and *semnum*. The level of permission required for each operation  
 39468 is shown with each command; see Section 2.7 (on page 489). The symbolic names for the values  
 39469 of *cmd* are defined in the `<sys/sem.h>` header:

39470 GETVAL Return the value of *semval*; see `<sys/sem.h>`. Requires read permission.

39471 SETVAL Set the value of *semval* to *arg.val*, where *arg* is the value of the fourth argument  
 39472 to *semctl()*. When this command is successfully executed, the *semadj* value  
 39473 corresponding to the specified semaphore in all processes is cleared. Requires  
 39474 alter permission; see Section 2.7 (on page 489).

39475 GETPID Return the value of *sempid*. Requires read permission.

39476 GETNCNT Return the value of *semmcnt*. Requires read permission.

39477 GETZCNT Return the value of *semzcnt*. Requires read permission.

39478 The following values of *cmd* operate on each *semval* in the set of semaphores:

39479 GETALL Return the value of *semval* for each semaphore in the semaphore set and place  
 39480 into the array pointed to by *arg.array*, where *arg* is the fourth argument to  
 39481 *semctl()*. Requires read permission.

39482 SETALL Set the value of *semval* for each semaphore in the semaphore set according to  
 39483 the array pointed to by *arg.array*, where *arg* is the fourth argument to *semctl()*.  
 39484 When this command is successfully executed, the *semadj* values corresponding  
 39485 to each specified semaphore in all processes are cleared. Requires alter  
 39486 permission.

39487 The following values of *cmd* are also available:

39488 IPC\_STAT Place the current value of each member of the **semid\_ds** data structure  
 39489 associated with *semid* into the structure pointed to by *arg.buf*, where *arg* is the  
 39490 fourth argument to *semctl()*. The contents of this structure are defined in

39491		<sys/sem.h>. Requires read permission.
39492	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> :
39493		
39494		
39495		<i>sem_perm.uid</i>
39496		<i>sem_perm.gid</i>
39497		<i>sem_perm.mode</i>
39498		The mode bits specified in Section 2.7.1 (on page 490) 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.
39499		
39500		
39501		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> .
39502		
39503		
39504		
39505	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> .
39506		
39507		
39508		
39509		
39510		
39511	<b>RETURN VALUE</b>	
39512	If successful, the value returned by <i>semctl()</i> depends on <i>cmd</i> as follows:	
39513	GETVAL	The value of <i>semval</i> .
39514	GETPID	The value of <i>sempid</i> .
39515	GETNCNT	The value of <i>semmcnt</i> .
39516	GETZCNT	The value of <i>semzcnt</i> .
39517	All others	0.
39518	Otherwise, <i>semctl()</i> shall return $-1$ and set <i>errno</i> to indicate the error.	
39519	<b>ERRORS</b>	
39520	The <i>semctl()</i> function shall fail if:	
39521	[EACCES]	Operation permission is denied to the calling process; see Section 2.7 (on page 489).
39522		
39523	[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.
39524		
39525		
39526	[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> .
39527		
39528		
39529		
39530	[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.
39531		

39532 **EXAMPLES**

39533 None.

39534 **APPLICATION USAGE**

39535 The fourth parameter in the SYNOPSIS section is now specified as ". . ." in order to avoid a  
39536 clash with the ISO C standard when referring to the union *semun* (as defined in Issue 3) and for  
39537 backward compatibility.

39538 The POSIX Realtime Extension defines alternative interfaces for interprocess communication.  
39539 Application developers who need to use IPC should design their applications so that modules  
39540 using the IPC routines described in Section 2.7 (on page 489) can be easily modified to use the  
39541 alternative interfaces.

39542 **RATIONALE**

39543 None.

39544 **FUTURE DIRECTIONS**

39545 None.

39546 **SEE ALSO**

39547 *semget()*, *semop()*, *sem\_close()*, *sem\_destroy()*, *sem\_getvalue()*, *sem\_init()*, *sem\_open()*, *sem\_post()*,  
39548 *sem\_unlink()*, *sem\_wait()*, the Base Definitions volume of IEEE Std 1003.1-200x, `<sys/sem.h>`,  
39549 Section 2.7 (on page 489)

39550 **CHANGE HISTORY**

39551 First released in Issue 2. Derived from Issue 2 of the SVID.

39552 **Issue 5**

39553 The note about use of POSIX Realtime Extension IPC routines has been moved from FUTURE  
39554 DIRECTIONS to the APPLICATION USAGE section.



39555 **NAME**39556 `semget` — get set of XSI semaphores39557 **SYNOPSIS**39558 XSI `#include <sys/sem.h>`39559 `int semget(key_t key, int nsems, int semflg);`

39560

39561 **DESCRIPTION**

39562 The `semget()` function operates on XSI semaphores (see the Base Definitions volume of |  
 39563 IEEE Std 1003.1-200x, Section 4.15, Semaphore). It is unspecified whether this function |  
 39564 interoperates with the realtime interprocess communication facilities defined in Section 2.8 (on |  
 39565 page 491).

39566 The `semget()` function shall return the semaphore identifier associated with *key*.

39567 A semaphore identifier with its associated `semid_ds` data structure and its associated set of |  
 39568 *nsems* semaphores (see `<sys/sem.h>`) is created for *key* if one of the following is true:

- 39569 • The argument *key* is equal to `IPC_PRIVATE`.
- 39570 • The argument *key* does not already have a semaphore identifier associated with it and (*semflg* |  
 39571 `&IPC_CREAT`) is non-zero.

39572 Upon creation, the `semid_ds` data structure associated with the new semaphore identifier is |  
 39573 initialized as follows:

- 39574 • In the operation permissions structure `sem_perm.cuid`, `sem_perm.uid`, `sem_perm.cgid`, and |  
 39575 `sem_perm.gid` shall be set equal to the effective user ID and effective group ID, respectively, of |  
 39576 the calling process.
- 39577 • The low-order 9 bits of `sem_perm.mode` shall be set equal to the low-order 9 bits of *semflg*. |
- 39578 • The variable `sem_nsems` shall be set equal to the value of *nsems*. |
- 39579 • The variable `sem_otime` shall be set equal to 0 and `sem_ctime` shall be set equal to the current |  
 39580 time.
- 39581 • The data structure associated with each semaphore in the set shall not be initialized. The |  
 39582 `semctl()` function with the command `SETVAL` or `SETALL` can be used to initialize each |  
 39583 semaphore.

39584 **RETURN VALUE**

39585 Upon successful completion, `semget()` shall return a non-negative integer, namely a semaphore |  
 39586 identifier; otherwise, it shall return `-1` and set `errno` to indicate the error.

39587 **ERRORS**

39588 The `semget()` function shall fail if:

- 39589 [EACCES] A semaphore identifier exists for *key*, but operation permission as specified by |  
 39590 the low-order 9 bits of *semflg* would not be granted; see Section 2.7 (on page |  
 39591 489).
- 39592 [EEXIST] A semaphore identifier exists for the argument *key* but ((*semflg* `&IPC_CREAT`) |  
 39593 `&&(semflg &IPC_EXCL)`) is non-zero.
- 39594 [EINVAL] The value of *nsems* is either less than or equal to 0 or greater than the system- |  
 39595 imposed limit, or a semaphore identifier exists for the argument *key*, but the |  
 39596 number of semaphores in the set associated with it is less than *nsems* and |  
 39597 *nsems* is not equal to 0.

39598 [ENOENT] A semaphore identifier does not exist for the argument *key* and (*semflg*  
39599 &IPC\_CREAT) is equal to 0.

39600 [ENOSPC] A semaphore identifier is to be created but the system-imposed limit on the  
39601 maximum number of allowed semaphores system-wide would be exceeded.

39602 **EXAMPLES**39603 **Creating a Semaphore Identifier**

39604 The following example gets a unique semaphore key using the *ftok()* function, then gets a  
39605 semaphore ID associated with that key using the *semget()* function (the first call also tests to  
39606 make sure the semaphore exists). If the semaphore does not exist, the program creates it, as  
39607 shown by the second call to *semget()*. In creating the semaphore for the queuing process, the  
39608 program attempts to create one semaphore with read/write permission for all. It also uses the  
39609 IPC\_EXCL flag, which forces *semget()* to fail if the semaphore already exists.

39610 After creating the semaphore, the program uses a call to *semop()* to initialize it to the values in  
39611 the *sbuf* array. The number of processes that can execute concurrently without queuing is  
39612 initially set to 2. The final call to *semget()* creates a semaphore identifier that can be used later in  
39613 the program.

```

39614 #include <sys/types.h>
39615 #include <stdio.h>
39616 #include <sys/ipc.h>
39617 #include <sys/sem.h>
39618 #include <sys/stat.h>
39619 #include <errno.h>
39620 #include <unistd.h>
39621 #include <stdlib.h>
39622 #include <pwd.h>
39623 #include <fcntl.h>
39624 #include <limits.h>
39625 ...
39626 key_t semkey;
39627 int semid, pfd, fv;
39628 struct sembuf sbuf;
39629 char *lgn;
39630 char filename[PATH_MAX+1];
39631 struct stat outstat;
39632 struct passwd *pw;
39633 ...
39634 /* Get unique key for semaphore. */
39635 if ((semkey = ftok("/tmp", 'a')) == (key_t) -1) {
39636     perror("IPC error: ftok"); exit(1);
39637 }
39638 /* Get semaphore ID associated with this key. */
39639 if ((semid = semget(semkey, 0, 0)) == -1) {
39640     /* Semaphore does not exist - Create. */
39641     if ((semid = semget(semkey, 1, IPC_CREAT | IPC_EXCL | S_IRUSR |
39642         S_IWUSR | S_IRGRP | S_IWGRP | S_IROTH | S_IWOTH)) != -1)
39643     {
39644         /* Initialize the semaphore. */
39645         sbuf.sem_num = 0;

```

```
39646         sbuf.sem_op = 2; /* This is the number of runs without queuing. */
39647         sbuf.sem_flg = 0;
39648         if (semop(semid, &sbuf, 1) == -1) {
39649             perror("IPC error: semop"); exit(1);
39650         }
39651     }
39652     else if (errno == EEXIST) {
39653         if ((semid = semget(semkey, 0, 0)) == -1) {
39654             perror("IPC error 1: semget"); exit(1);
39655         }
39656     }
39657     else {
39658         perror("IPC error 2: semget"); exit(1);
39659     }
39660 }
39661 ...
```

#### 39662 APPLICATION USAGE

39663 The POSIX Realtime Extension defines alternative interfaces for interprocess communication.  
39664 Application developers who need to use IPC should design their applications so that modules  
39665 using the IPC routines described in Section 2.7 (on page 489) can be easily modified to use the  
39666 alternative interfaces.

#### 39667 RATIONALE

39668 None.

#### 39669 FUTURE DIRECTIONS

39670 None.

#### 39671 SEE ALSO

39672 *semctl()*, *semop()*, *sem\_close()*, *sem\_destroy()*, *sem\_getvalue()*, *sem\_init()*, *sem\_open()*, *sem\_post()*,  
39673 *sem\_unlink()*, *sem\_wait()*, the Base Definitions volume of IEEE Std 1003.1-200x, <sys/sem.h>,  
39674 Section 2.7 (on page 489).

#### 39675 CHANGE HISTORY

39676 First released in Issue 2. Derived from Issue 2 of the SVID.

#### 39677 Issue 5

39678 The note about use of POSIX Realtime Extension IPC routines has been moved from FUTURE  
39679 DIRECTIONS to a new APPLICATION USAGE section.

## 39680 NAME

39681 semop — XSI semaphore operations

## 39682 SYNOPSIS

39683 XSI #include &lt;sys/sem.h&gt;

39684 int semop(int *semid*, struct sembuf \**sops*, size\_t *nsops*);

39685

## 39686 DESCRIPTION

39687 The *semop()* function operates on XSI semaphores (see the Base Definitions volume of |  
 39688 IEEE Std 1003.1-200x, Section 4.15, Semaphore). It is unspecified whether this function |  
 39689 interoperates with the realtime interprocess communication facilities defined in Section 2.8 (on |  
 39690 page 491).

39691 The *semop()* function shall perform atomically a user-defined array of semaphore operations on |  
 39692 the set of semaphores associated with the semaphore identifier specified by the argument *semid*. |

39693 The argument *sops* is a pointer to a user-defined array of semaphore operation structures. The |  
 39694 implementation shall not modify elements of this array unless the application uses |  
 39695 implementation-defined extensions.

39696 The argument *nsops* is the number of such structures in the array.

39697 Each structure, **sembuf**, includes the following members:

39698

39699

39700

39701

39702

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.

39703 Each semaphore operation specified by *sem\_op* is performed on the corresponding semaphore |  
 39704 specified by *semid* and *sem\_num*.

39705 The variable *sem\_op* specifies one of three semaphore operations:

39706 1. If *sem\_op* is a negative integer and the calling process has alter permission, one of the |  
 39707 following shall occur:

39708 • If *semval* (see <sys/sem.h>) is greater than or equal to the absolute value of *sem\_op*, the |  
 39709 absolute value of *sem\_op* is subtracted from *semval*. Also, if (*sem\_flg* &SEM\_UNDO) is |  
 39710 non-zero, the absolute value of *sem\_op* shall be added to the calling process' *semadj* |  
 39711 value for the specified semaphore.

39712 • If *semval* is less than the absolute value of *sem\_op* and (*sem\_flg* &IPC\_NOWAIT) is non- |  
 39713 zero, *semop()* shall return immediately.

39714 • If *semval* is less than the absolute value of *sem\_op* and (*sem\_flg* &IPC\_NOWAIT) is 0, |  
 39715 *semop()* shall increment the *semncnt* associated with the specified semaphore and |  
 39716 suspend execution of the calling thread until one of the following conditions occurs:

39717 — The value of *semval* becomes greater than or equal to the absolute value of *sem\_op*. |  
 39718 When this occurs, the value of *semncnt* associated with the specified semaphore |  
 39719 shall be decremented, the absolute value of *sem\_op* shall be subtracted from *semval* |  
 39720 and, if (*sem\_flg* &SEM\_UNDO) is non-zero, the absolute value of *sem\_op* shall be |  
 39721 added to the calling process' *semadj* value for the specified semaphore. |

39722 — The *semid* for which the calling thread is awaiting action is removed from the |  
 39723 system. When this occurs, *errno* shall be set equal to [EIDRM] and *-1* shall be

39724 returned.

39725 — The calling thread receives a signal that is to be caught. When this occurs, the value  
39726 of *semncnt* associated with the specified semaphore shall be decremented, and the  
39727 calling thread shall resume execution in the manner prescribed in *sigaction()*.

39728 2. If *sem\_op* is a positive integer and the calling process has alter permission, the value of  
39729 *sem\_op* shall be added to *semval* and, if (*sem\_flg* &SEM\_UNDO) is non-zero, the value of  
39730 *sem\_op* shall be subtracted from the calling process' *semadj* value for the specified  
39731 semaphore.

39732 3. If *sem\_op* is 0 and the calling process has read permission, one of the following shall occur:

39733 • If *semval* is 0, *semop()* shall return immediately.

39734 • If *semval* is non-zero and (*sem\_flg* &IPC\_NOWAIT) is non-zero, *semop()* shall return  
39735 immediately.

39736 • If *semval* is non-zero and (*sem\_flg* &IPC\_NOWAIT) is 0, *semop()* shall increment the  
39737 *semzcnt* associated with the specified semaphore and suspend execution of the calling  
39738 thread until one of the following occurs:

39739 — The value of *semval* becomes 0, at which time the value of *semzcnt* associated with  
39740 the specified semaphore shall be decremented.

39741 — The *semid* for which the calling thread is awaiting action is removed from the  
39742 system. When this occurs, *errno* shall be set equal to [EIDRM] and -1 shall be  
39743 returned.

39744 — The calling thread receives a signal that is to be caught. When this occurs, the value  
39745 of *semzcnt* associated with the specified semaphore shall be decremented, and the  
39746 calling thread shall resume execution in the manner prescribed in *sigaction()*.

39747 Upon successful completion, the value of *sempid* for each semaphore specified in the array  
39748 pointed to by *sops* shall be set equal to the process ID of the calling process.

39749 **RETURN VALUE**

39750 Upon successful completion, *semop()* shall return 0; otherwise, it shall return -1 and set *errno* to  
39751 indicate the error.

39752 **ERRORS**

39753 The *semop()* function shall fail if:

39754 [E2BIG] The value of *nsops* is greater than the system-imposed maximum.

39755 [EACCES] Operation permission is denied to the calling process; see Section 2.7 (on page  
39756 489).

39757 [EAGAIN] The operation would result in suspension of the calling process but (*sem\_flg*  
39758 &IPC\_NOWAIT) is non-zero.

39759 [EFBIG] The value of *sem\_num* is less than 0 or greater than or equal to the number of  
39760 semaphores in the set associated with *semid*.

39761 [EIDRM] The semaphore identifier *semid* is removed from the system.

39762 [EINTR] The *semop()* function was interrupted by a signal.

39763 [EINVAL] The value of *semid* is not a valid semaphore identifier, or the number of  
39764 individual semaphores for which the calling process requests a SEM\_UNDO  
39765 would exceed the system-imposed limit.

39766 [ENOSPC] The limit on the number of individual processes requesting a SEM\_UNDO  
 39767 would be exceeded.

39768 [ERANGE] An operation would cause a *semval* to overflow the system-imposed limit, or  
 39769 an operation would cause a *semadj* value to overflow the system-imposed  
 39770 limit.

### 39771 EXAMPLES

#### 39772 Setting Values in Semaphores

39773 The following example sets the values of the two semaphores associated with the *semid*  
 39774 identifier to the values contained in the *sb* array.

```
39775 #include <sys/sem.h>
39776 ...
39777 int semid;
39778 struct sembuf sb[2];
39779 int nsops = 2;
39780 int result;

39781 /* Adjust value of semaphore in the semaphore array semid. */
39782 sb[0].sem_num = 0;
39783 sb[0].sem_op = -1;
39784 sb[0].sem_flg = SEM_UNDO | IPC_NOWAIT;
39785 sb[1].sem_num = 1;
39786 sb[1].sem_op = 1;
39787 sb[1].sem_flg = 0;

39788 result = semop(semid, sb, nsops);
```

#### 39789 Creating a Semaphore Identifier

39790 The following example gets a unique semaphore key using the *ftok()* function, then gets a  
 39791 semaphore ID associated with that key using the *semget()* function (the first call also tests to  
 39792 make sure the semaphore exists). If the semaphore does not exist, the program creates it, as  
 39793 shown by the second call to *semget()*. In creating the semaphore for the queuing process, the  
 39794 program attempts to create one semaphore with read/write permission for all. It also uses the  
 39795 IPC\_EXCL flag, which forces *semget()* to fail if the semaphore already exists.

39796 After creating the semaphore, the program uses a call to *semop()* to initialize it to the values in  
 39797 the *sbuf* array. The number of processes that can execute concurrently without queuing is  
 39798 initially set to 2. The final call to *semget()* creates a semaphore identifier that can be used later in  
 39799 the program.

39800 The final call to *semop()* acquires the semaphore and waits until it is free; the SEM\_UNDO  
 39801 option releases the semaphore when the process exits, waiting until there are less than two  
 39802 processes running concurrently.

```
39803 #include <sys/types.h>
39804 #include <stdio.h>
39805 #include <sys/ipc.h>
39806 #include <sys/sem.h>
39807 #include <sys/stat.h>
39808 #include <errno.h>
39809 #include <unistd.h>
39810 #include <stdlib.h>
```

```

39811     #include <pwd.h>
39812     #include <fcntl.h>
39813     #include <limits.h>
39814     ...
39815     key_t semkey;
39816     int semid, pfd, fv;
39817     struct sembuf sbuf;
39818     char *lgn;
39819     char filename[PATH_MAX+1];
39820     struct stat outstat;
39821     struct passwd *pw;
39822     ...
39823     /* Get unique key for semaphore. */
39824     if ((semkey = ftok("/tmp", 'a')) == (key_t) -1) {
39825         perror("IPC error: ftok"); exit(1);
39826     }
39827     /* Get semaphore ID associated with this key. */
39828     if ((semid = semget(semkey, 0, 0)) == -1) {
39829         /* Semaphore does not exist - Create. */
39830         if ((semid = semget(semkey, 1, IPC_CREAT | IPC_EXCL | S_IRUSR |
39831             S_IWUSR | S_IRGRP | S_IWGRP | S_IROTH | S_IWOTH)) != -1)
39832         {
39833             /* Initialize the semaphore. */
39834             sbuf.sem_num = 0;
39835             sbuf.sem_op = 2; /* This is the number of runs without queuing. */
39836             sbuf.sem_flg = 0;
39837             if (semop(semid, &sbuf, 1) == -1) {
39838                 perror("IPC error: semop"); exit(1);
39839             }
39840         }
39841         else if (errno == EEXIST) {
39842             if ((semid = semget(semkey, 0, 0)) == -1) {
39843                 perror("IPC error 1: semget"); exit(1);
39844             }
39845         }
39846         else {
39847             perror("IPC error 2: semget"); exit(1);
39848         }
39849     }
39850     ...
39851     sbuf.sem_num = 0;
39852     sbuf.sem_op = -1;
39853     sbuf.sem_flg = SEM_UNDO;
39854     if (semop(semid, &sbuf, 1) == -1) {
39855         perror("IPC Error: semop"); exit(1);
39856     }

```

### 39857 APPLICATION USAGE

39858 The POSIX Realtime Extension defines alternative interfaces for interprocess communication.  
39859 Application developers who need to use IPC should design their applications so that modules  
39860 using the IPC routines described in Section 2.7 (on page 489) can be easily modified to use the  
39861 alternative interfaces.

39862 **RATIONALE**

39863           None.

39864 **FUTURE DIRECTIONS**

39865           None.

39866 **SEE ALSO**

39867           *exec*, *exit()*, *fork()*, *semctl()*, *semget()*, *sem\_close()*, *sem\_destroy()*, *sem\_getvalue()*, *sem\_init()*,  
39868           *sem\_open()*, *sem\_post()*, *sem\_unlink()*, *sem\_wait()*, the Base Definitions volume of  
39869           IEEE Std 1003.1-200x, <sys/ipc.h>, <sys/sem.h>, <sys/types.h>, Section 2.7 (on page 489)

39870 **CHANGE HISTORY**

39871           First released in Issue 2. Derived from Issue 2 of the SVID.

39872 **Issue 5**

39873           The note about use of POSIX Realtime Extension IPC routines has been moved from FUTURE  
39874           DIRECTIONS to a new APPLICATION USAGE section.



39875 **NAME**

39876           send — send a message on a socket

39877 **SYNOPSIS**

39878           #include &lt;sys/socket.h&gt;

39879           ssize\_t send(int *socket*, const void \**buffer*, size\_t *length*, int *flags*);39880 **DESCRIPTION**39881           The *send()* function shall initiate transmission of a message from the specified socket to its peer.39882           The *send()* function shall send a message only when the socket is connected (including when the peer of a connectionless socket has been set via *connect()*).39884           The *send()* functions takes the following arguments:39885           *socket*               Specifies the socket file descriptor.39886           *buffer*               Points to the buffer containing the message to send.39887           *length*               Specifies the length of the message in bytes.39888           *flags*               Specifies the type of message transmission. Values of this argument are formed by logically OR'ing zero or more of the following flags:

39890                           MSG\_EOR       Terminates a record (if supported by the protocol).

39891                           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.

39894           The length of the message to be sent is specified by the *length* argument. If the message is too long to pass through the underlying protocol, *send()* shall fail and no data shall be transmitted.39896           Successful completion of a call to *send()* does not guarantee delivery of the message. A return value of  $-1$  indicates only locally-detected errors.39898           If space is not available at the sending socket to hold the message to be transmitted, and the socket file descriptor does not have `O_NONBLOCK` set, *send()* shall block until space is available. If space is not available at the sending socket to hold the message to be transmitted, and the socket file descriptor does have `O_NONBLOCK` set, *send()* shall fail. The *select()* and *poll()* functions can be used to determine when it is possible to send more data.39903           The socket in use may require the process to have appropriate privileges to use the *send()* function.39905 **RETURN VALUE**39906           Upon successful completion, *send()* shall return the number of bytes sent. Otherwise,  $-1$  shall be returned and *errno* set to indicate the error.39908 **ERRORS**39909           The *send()* function shall fail if:

39910           [EAGAIN] or [EWOULDBLOCK]

39911                           The socket's file descriptor is marked `O_NONBLOCK` and the requested operation would block.39913           [EBADF]           The *socket* argument is not a valid file descriptor.

39914           [ECONNRESET] A connection was forcibly closed by a peer.

39915           [EDESTADDRREQ]

39916                           The socket is not connection-mode and no peer address is set.

- 39917 [EINTR] A signal interrupted *send()* before any data was transmitted.
- 39918 [EMSGSIZE] The message is too large to be sent all at once, as the socket requires.
- 39919 [ENOTCONN] The socket is not connected or otherwise has not had the peer pre-specified.
- 39920 [ENOTSOCK] The *socket* argument does not refer to a socket.
- 39921 [EOPNOTSUPP] The *socket* argument is associated with a socket that does not support one or  
39922 more of the values set in *flags*.
- 39923 [EPIPE] The socket is shut down for writing, or the socket is connection-mode and is  
39924 no longer connected. In the latter case, and if the socket is of type  
39925 SOCK\_STREAM, the SIGPIPE signal is generated to the calling thread.
- 39926 The *send()* function may fail if:
- 39927 [EACCES] The calling process does not have the appropriate privileges.
- 39928 [EIO] An I/O error occurred while reading from or writing to the file system.
- 39929 [ENETDOWN] The local network interface used to reach the destination is down.
- 39930 [ENETUNREACH]  
39931 No route to the network is present.
- 39932 [ENOBUFS] Insufficient resources were available in the system to perform the operation.
- 39933 **EXAMPLES**
- 39934 None.
- 39935 **APPLICATION USAGE**
- 39936 The *send()* function is equivalent to *sendto()* with a null pointer *dest\_len* argument, and to *write()* |  
39937 if no flags are used.
- 39938 **RATIONALE**
- 39939 None.
- 39940 **FUTURE DIRECTIONS**
- 39941 None.
- 39942 **SEE ALSO**
- 39943 *connect()*, *getsockopt()*, *poll()*, *recv()*, *recvfrom()*, *recvmsg()*, *select()*, *sendmsg()*, *sendto()*,  
39944 *setsockopt()*, *shutdown()*, *socket()*, the Base Definitions volume of IEEE Std 1003.1-200x,  
39945 <sys/socket.h>
- 39946 **CHANGE HISTORY**
- 39947 First released in Issue 6. Derived from the XNS, Issue 5.2 specification.

39948 **NAME**

39949 sendmsg — send a message on a socket using a message structure

39950 **SYNOPSIS**

39951 #include &lt;sys/socket.h&gt;

39952 ssize\_t sendmsg(int *socket*, const struct msghdr \**message*, int *flags*);39953 **DESCRIPTION**

39954 The *sendmsg()* function shall send a message through a connection-mode or connectionless-mode socket. If the socket is connectionless-mode, the message shall be sent to the address specified by **msghdr**. If the socket is connection-mode, the destination address in **msghdr** shall be ignored.

39958 The *sendmsg()* function takes the following arguments:

39959	<i>socket</i>	Specifies the socket file descriptor.
39960	<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.
39963	<i>flags</i>	Specifies the type of message transmission. The application may specify 0 or the following flag:
39965	MSG_EOR	Terminates a record (if supported by the protocol).
39966	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.

39969 The *msg\_iov* and *msg\_iovlen* fields of *message* specify zero or more buffers containing the data to be sent. *msg\_iov* points to an array of **iovec** structures; *msg\_iovlen* shall be set to the dimension of this array. In each **iovec** structure, the *iov\_base* field specifies a storage area and the *iov\_len* field gives its size in bytes. Some of these sizes can be zero. The data from each storage area indicated by *msg\_iov* is sent in turn.

39974 Successful completion of a call to *sendmsg()* does not guarantee delivery of the message. A return value of  $-1$  indicates only locally-detected errors.

39976 If space is not available at the sending socket to hold the message to be transmitted and the socket file descriptor does not have **O\_NONBLOCK** set, *sendmsg()* function shall block until space is available. If space is not available at the sending socket to hold the message to be transmitted and the socket file descriptor does have **O\_NONBLOCK** set, *sendmsg()* function shall fail.

39981 If the socket protocol supports broadcast and the specified address is a broadcast address for the socket protocol, *sendmsg()* shall fail if the **SO\_BROADCAST** option is not set for the socket.

39983 The socket in use may require the process to have appropriate privileges to use the *sendmsg()* function.

39985 **RETURN VALUE**

39986 Upon successful completion, *sendmsg()* shall return the number of bytes sent. Otherwise,  $-1$  shall be returned and *errno* set to indicate the error.

39988 **ERRORS**39989 The *sendmsg()* function shall fail if:

39990 [EAGAIN] or [EWOULDBLOCK]

39991 The socket's file descriptor is marked **O\_NONBLOCK** and the requested

39992		operation would block.
39993	[EAFNOSUPPORT]	
39994		Addresses in the specified address family cannot be used with this socket.
39995	[EBADF]	The <i>socket</i> argument is not a valid file descriptor.
39996	[ECONNRESET]	A connection was forcibly closed by a peer.
39997	[EINTR]	A signal interrupted <i>sendmsg()</i> before any data was transmitted.
39998	[EINVAL]	The sum of the <i>iov_len</i> values overflows an <i>ssize_t</i> .
39999	[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 <i>msghdr</i> structure pointed to by <i>message</i> is less than or equal to 0 or is greater than {IOV_MAX}.
40000		
40001		
40002	[ENOTCONN]	The socket is connection-mode but is not connected.
40003	[ENOTSOCK]	The <i>socket</i> argument does not refer a socket.
40004	[EOPNOTSUPP]	The <i>socket</i> argument is associated with a socket that does not support one or more of the values set in <i>flags</i> .
40005		
40006	[EPIPE]	The socket is shut down for writing, or the socket is connection-mode and is no longer connected. In the latter case, and if the socket is of type SOCK_STREAM, the SIGPIPE signal is generated to the calling thread.
40007		
40008		
40009		If the address family of the socket is AF_UNIX, then <i>sendmsg()</i> shall fail if:
40010	[EIO]	An I/O error occurred while reading from or writing to the file system.
40011	[ELOOP]	A loop exists in symbolic links encountered during resolution of the pathname
40012		in the socket address.
40013	[ENAMETOOLONG]	
40014		A component of a pathname exceeded {NAME_MAX} characters, or an entire
40015		pathname exceeded {PATH_MAX} characters.
40016	[ENOENT]	A component of the pathname does not name an existing file or the path name
40017		is an empty string.
40018	[ENOTDIR]	A component of the path prefix of the pathname in the socket address is not a
40019		directory.
40020		The <i>sendmsg()</i> function may fail if:
40021	[EACCES]	Search permission is denied for a component of the path prefix; or write
40022		access to the named socket is denied.
40023	[EDESTADDRREQ]	
40024		The socket is not connection-mode and does not have its peer address set, and
40025		no destination address was specified.
40026	[EHOSTUNREACH]	
40027		The destination host cannot be reached (probably because the host is down or
40028		a remote router cannot reach it).
40029	[EIO]	An I/O error occurred while reading from or writing to the file system.
40030	[EISCONN]	A destination address was specified and the socket is already connected.
40031	[ENETDOWN]	The local network interface used to reach the destination is down.

- 40032 [ENETUNREACH]  
40033 No route to the network is present.
- 40034 [ENOBUFS] Insufficient resources were available in the system to perform the operation.
- 40035 [ENOMEM] Insufficient memory was available to fulfill the request.
- 40036 If the address family of the socket is AF\_UNIX, then *sendmsg()* may fail if:
- 40037 [ELOOP] More than {SYMLOOP\_MAX} symbolic links were encountered during |  
40038 resolution of the pathname in the socket address. |
- 40039 [ENAMETOOLONG]  
40040 Pathname resolution of a symbolic link produced an intermediate result |  
40041 whose length exceeds {PATH\_MAX}. |
- 40042 **EXAMPLES**  
40043 Done.
- 40044 **APPLICATION USAGE**  
40045 The *select()* and *poll()* functions can be used to determine when it is possible to send more data.
- 40046 **RATIONALE**  
40047 None.
- 40048 **FUTURE DIRECTIONS**  
40049 None.
- 40050 **SEE ALSO**  
40051 *getsockopt()*, *poll()*, *recv()*, *recvfrom()*, *recvmsg()*, *select()*, *send()*, *sendto()*, *setsockopt()*,  
40052 *shutdown()*, *socket()*, the Base Definitions volume of IEEE Std 1003.1-200x, <sys/socket.h>
- 40053 **CHANGE HISTORY**  
40054 First released in Issue 6. Derived from the XNS, Issue 5.2 specification.
- 40055 The wording of the mandatory [ELOOP] error condition is updated, and a second optional  
40056 [ELOOP] error condition is added.

## 40057 NAME

40058 sendto — send a message on a socket

## 40059 SYNOPSIS

40060 #include &lt;sys/socket.h&gt;

40061 ssize\_t sendto(int socket, const void \*message, size\_t length,

40062 int flags, const struct sockaddr \*dest\_addr,

40063 socklen\_t dest\_len);

## 40064 DESCRIPTION

40065 The *sendto()* function shall send a message through a connection-mode or connectionless-mode  
 40066 socket. If the socket is connectionless-mode, the message shall be sent to the address specified by  
 40067 *dest\_addr*. If the socket is connection-mode, *dest\_addr* shall be ignored.

40068 The *sendto()* function takes the following arguments:

40069	<i>socket</i>	Specifies the socket file descriptor.
40070	<i>message</i>	Points to a buffer containing the message to be sent.
40071	<i>length</i>	Specifies the size of the message in bytes.
40072	<i>flags</i>	Specifies the type of message transmission. Values of this argument are 40073 formed by logically OR'ing zero or more of the following flags:
40074	MSG_EOR	Terminates a record (if supported by the protocol).
40075	MSG_OOB	Sends out-of-band data on sockets that support out-of-band 40076 data. The significance and semantics of out-of-band data are 40077 protocol-specific.
40078	<i>dest_addr</i>	Points to a <b>sockaddr</b> structure containing the destination address. The length 40079 and format of the address depend on the address family of the socket.
40080	<i>dest_len</i>	Specifies the length of the <b>sockaddr</b> structure pointed to by the <i>dest_addr</i> 40081 argument.

40082 If the socket protocol supports broadcast and the specified address is a broadcast address for the  
 40083 socket protocol, *sendto()* shall fail if the SO\_BROADCAST option is not set for the socket.

40084 The *dest\_addr* argument specifies the address of the target. The *length* argument specifies the  
 40085 length of the message.

40086 Successful completion of a call to *sendto()* does not guarantee delivery of the message. A return  
 40087 value of  $-1$  indicates only locally-detected errors.

40088 If space is not available at the sending socket to hold the message to be transmitted and the  
 40089 socket file descriptor does not have O\_NONBLOCK set, *sendto()* shall block until space is  
 40090 available. If space is not available at the sending socket to hold the message to be transmitted  
 40091 and the socket file descriptor does have O\_NONBLOCK set, *sendto()* shall fail.

40092 The socket in use may require the process to have appropriate privileges to use the *sendto()*  
 40093 function.

## 40094 RETURN VALUE

40095 Upon successful completion, *sendto()* shall return the number of bytes sent. Otherwise,  $-1$  shall  
 40096 be returned and *errno* set to indicate the error.

40097 **ERRORS**

- 40098       The *sendto()* function shall fail if:
- 40099       [EAFNOSUPPORT]
- 40100               Addresses in the specified address family cannot be used with this socket.
- 40101       [EAGAIN] or [EWOULDBLOCK]
- 40102               The socket's file descriptor is marked O\_NONBLOCK and the requested
- 40103               operation would block.
- 40104       [EBADF]       The *socket* argument is not a valid file descriptor.
- 40105       [ECONNRESET] A connection was forcibly closed by a peer.
- 40106       [EINTR]       A signal interrupted *sendto()* before any data was transmitted.
- 40107       [EMSGSIZE]   The message is too large to be sent all at once, as the socket requires.
- 40108       [ENOTCONN]   The socket is connection-mode but is not connected.
- 40109       [ENOTSOCK]   The *socket* argument does not refer to a socket.
- 40110       [EOPNOTSUPP] The *socket* argument is associated with a socket that does not support one or
- 40111               more of the values set in *flags*.
- 40112       [EPIPE]       The socket is shut down for writing, or the socket is connection-mode and is
- 40113               no longer connected. In the latter case, and if the socket is of type
- 40114               SOCK\_STREAM, the SIGPIPE signal is generated to the calling thread.
- 40115       If the address family of the socket is AF\_UNIX, then *sendto()* shall fail if:
- 40116       [EIO]        An I/O error occurred while reading from or writing to the file system.
- 40117       [ELOOP]       A loop exists in symbolic links encountered during resolution of the pathname |
- 40118               in the socket address. |
- 40119       [ENAMETOOLONG]
- 40120               A component of a pathname exceeded {NAME\_MAX} characters, or an entire |
- 40121               pathname exceeded {PATH\_MAX} characters. |
- 40122       [ENOENT]       A component of the pathname does not name an existing file or the pathname |
- 40123               is an empty string. |
- 40124       [ENOTDIR]     A component of the path prefix of the pathname in the socket address is not a |
- 40125               directory. |
- 40126       The *sendto()* function may fail if:
- 40127       [EACCES]       Search permission is denied for a component of the path prefix; or write
- 40128               access to the named socket is denied.
- 40129       [EDESTADDRREQ]
- 40130               The socket is not connection-mode and does not have its peer address set, and
- 40131               no destination address was specified.
- 40132       [EHOSTUNREACH]
- 40133               The destination host cannot be reached (probably because the host is down or
- 40134               a remote router cannot reach it).
- 40135       [EINVAL]       The *dest\_len* argument is not a valid length for the address family.
- 40136       [EIO]        An I/O error occurred while reading from or writing to the file system.

- 40137 [EISCONN] A destination address was specified and the socket is already connected. This  
40138 error may or may not be returned for connection mode sockets.
- 40139 [ENETDOWN] The local network interface used to reach the destination is down.
- 40140 [ENETUNREACH]  
40141 No route to the network is present.
- 40142 [ENOBUFS] Insufficient resources were available in the system to perform the operation.
- 40143 [ENOMEM] Insufficient memory was available to fulfill the request.
- 40144 If the address family of the socket is AF\_UNIX, then *sendto()* may fail if:
- 40145 [ELOOP] More than {SYMLOOP\_MAX} symbolic links were encountered during |  
40146 resolution of the pathname in the socket address. |
- 40147 [ENAMETOOLONG]  
40148 Pathname resolution of a symbolic link produced an intermediate result |  
40149 whose length exceeds {PATH\_MAX}.
- 40150 **EXAMPLES**  
40151 None.
- 40152 **APPLICATION USAGE**  
40153 The *select()* and *poll()* functions can be used to determine when it is possible to send more data.
- 40154 **RATIONALE**  
40155 None.
- 40156 **FUTURE DIRECTIONS**  
40157 None.
- 40158 **SEE ALSO**  
40159 *getsockopt()*, *poll()*, *recv()*, *recvfrom()*, *recvmsg()*, *select()*, *send()*, *sendmsg()*, *setsockopt()*,  
40160 *shutdown()*, *socket()*, the Base Definitions volume of IEEE Std 1003.1-200x, <sys/socket.h>
- 40161 **CHANGE HISTORY**  
40162 First released in Issue 6. Derived from the XNS, Issue 5.2 specification.
- 40163 The wording of the mandatory [ELOOP] error condition is updated, and a second optional  
40164 [ELOOP] error condition is added.



40165 **NAME**

40166 setbuf — assign buffering to a stream

40167 **SYNOPSIS**

40168 #include &lt;stdio.h&gt;

40169 void setbuf(FILE \*restrict stream, char \*restrict buf);

40170 **DESCRIPTION**

40171 cx The functionality described on this reference page is aligned with the ISO C standard. Any  
40172 conflict between the requirements described here and the ISO C standard is unintentional. This  
40173 volume of IEEE Std 1003.1-200x defers to the ISO C standard.

40174 Except that it returns no value, the function call:

40175 setbuf(stream, buf)

40176 shall be equivalent to:

40177 setvbuf(stream, buf, \_IOFBF, BUFSIZ)

40178 if *buf* is not a null pointer, or to:

40179 setvbuf(stream, buf, \_IONBF, BUFSIZ)

40180 if *buf* is a null pointer.40181 **RETURN VALUE**40182 The *setbuf()* function shall not return a value.40183 **ERRORS**

40184 No errors are defined.

40185 **EXAMPLES**

40186 None.

40187 **APPLICATION USAGE**

40188 A common source of error is allocating buffer space as an “automatic” variable in a code block,  
40189 and then failing to close the stream in the same block.

40190 With *setbuf()*, allocating a buffer of BUFSIZ bytes does not necessarily imply that all of BUFSIZ  
40191 bytes are used for the buffer area.

40192 **RATIONALE**

40193 None.

40194 **FUTURE DIRECTIONS**

40195 None.

40196 **SEE ALSO**40197 *fopen()*, *setvbuf()*, the Base Definitions volume of IEEE Std 1003.1-200x, <stdio.h>40198 **CHANGE HISTORY**

40199 First released in Issue 1. Derived from Issue 1 of the SVID.

40200 **Issue 6**40201 The prototype for *setbuf()* is updated for alignment with the ISO/IEC 9899:1999 standard.

40202 **NAME**

40203           setcontext — set current user context

40204 **SYNOPSIS**

40205 xSI       #include <ucontext.h>

40206           int setcontext(const ucontext\_t \*ucp);

40207

40208 **DESCRIPTION**

40209           Refer to *getcontext()*.

40210 **NAME**

40211 setegid — set effective group ID

40212 **SYNOPSIS**

40213 #include &lt;unistd.h&gt;

40214 int setegid(gid\_t gid);

40215 **DESCRIPTION**

40216 If *gid* is equal to the real group ID or the saved set-group-ID, or if the process has appropriate  
40217 privileges, *setegid()* shall set the effective group ID of the calling process to *gid*; the real group  
40218 ID, saved set-group-ID, and any supplementary group IDs shall remain unchanged.

40219 The *setegid()* function shall not affect the supplementary group list in any way.

40220 **RETURN VALUE**

40221 Upon successful completion, 0 shall be returned; otherwise, -1 shall be returned and *errno* set to  
40222 indicate the error.

40223 **ERRORS**

40224 The *setegid()* function shall fail if:

40225 [EINVAL] The value of the *gid* argument is invalid and is not supported by the  
40226 implementation.

40227 [EPERM] The process does not have appropriate privileges and *gid* does not match the  
40228 real group ID or the saved set-group-ID.

40229 **EXAMPLES**

40230 None.

40231 **APPLICATION USAGE**

40232 None.

40233 **RATIONALE**40234 Refer to the RATIONALE section in *setuid()*.40235 **FUTURE DIRECTIONS**

40236 None.

40237 **SEE ALSO**

40238 *exec*, *getegid()*, *geteuid()*, *getgid()*, *getuid()*, *seteuid()*, *setgid()*, *setregid()*, *setreuid()*, *setuid()*, the  
40239 Base Definitions volume of IEEE Std 1003.1-200x, <sys/types.h>, <unistd.h>

40240 **CHANGE HISTORY**

40241 First released in Issue 6. Derived from the IEEE P1003.1a draft standard.

40242 **NAME**

40243        setenv — add or change environment variable

40244 **SYNOPSIS**40245 `cx        #include <stdlib.h>`40246        `int setenv(const char *envname, const char *envval, int overwrite);`

40247

40248 **DESCRIPTION**40249        The *setenv()* function shall update or add a variable in the environment of the calling process.40250        The *envname* argument points to a string containing the name of an environment variable to be40251        added or altered. The environment variable shall be set to the value to which *envval* points. The40252        function shall fail if *envname* points to a string which contains an '=' character. If the40253        environment variable named by *envname* already exists and the value of *overwrite* is non-zero,

40254        the function shall return success and the environment shall be updated. If the environment

40255        variable named by *envname* already exists and the value of *overwrite* is zero, the function shall

40256        return success and the environment shall remain unchanged.

40257        If the application modifies *environ* or the pointers to which it points, the behavior of *setenv()* is40258        undefined. The *setenv()* function shall update the list of pointers to which *environ* points.40259        The strings described by *envname* and *envval* are copied by this function.40260        The *setenv()* function need not be reentrant. A function that is not required to be reentrant is not

40261        required to be thread-safe.

40262 **RETURN VALUE**40263        Upon successful completion, zero shall be returned. Otherwise, -1 shall be returned, *errno* set to

40264        indicate the error, and the environment shall be unchanged.

40265 **ERRORS**40266        The *setenv()* function shall fail if:40267        [EINVAL]        The *name* argument is a null pointer, points to an empty string, or points to a  
40268        string containing an '=' character.40269        [ENOMEM]        Insufficient memory was available to add a variable or its value to the  
40270        environment.40271 **EXAMPLES**

40272        None.

40273 **APPLICATION USAGE**

40274        None.

40275 **RATIONALE**40276        Unanticipated results may occur if *setenv()* changes the external variable *environ*. In particular,40277        if the optional *envp* argument to *main()* is present, it is not changed, and thus may point to an40278        obsolete copy of the environment (as may any other copy of *environ*). However, other than the

40279        aforementioned restriction, the developers of IEEE Std 1003.1-200x intended that the traditional

40280        method of walking through the environment by way of the *environ* pointer must be supported.40281        It was decided that *setenv()* should be required by this revision because it addresses a piece of

40282        missing functionality, and does not impose a significant burden on the implementor.

40283        There was considerable debate as to whether the System V *putenv()* function or the BSD *setenv()*40284        function should be required as a mandatory function. The *setenv()* function was chosen because40285        it permitted the implementation of *unsetenv()* function to delete environmental variables,40286        without specifying an additional interface. The *putenv()* function is available as an XSI

40287 extension.

40288 The standard developers considered requiring that *setenv()* indicate an error when a call to it  
40289 would result in exceeding {ARG\_MAX}. The requirement was rejected since the condition might  
40290 be temporary, with the application eventually reducing the environment size. The ultimate  
40291 success or failure depends on the size at the time of a call to *exec*, which returns an indication of  
40292 this error condition.

40293 **FUTURE DIRECTIONS**

40294 None.

40295 **SEE ALSO**

40296 *getenv()*, *unsetenv()*, the Base Definitions volume of IEEE Std 1003.1-200x, `<stdlib.h>`,  
40297 `<sys/types.h>`, `<unistd.h>`

40298 **CHANGE HISTORY**

40299 First released in Issue 6. Derived from the IEEE P1003.1a draft standard.

40300 **NAME**

40301           seteuid — set effective user ID

40302 **SYNOPSIS**

40303           #include &lt;unistd.h&gt;

40304           int seteuid(uid\_t uid);

40305 **DESCRIPTION**

40306           If *uid* is equal to the real user ID or the saved set-user-ID, or if the process has appropriate  
40307           privileges, *seteuid()* shall set the effective user ID of the calling process to *uid*; the real user ID  
40308           and saved set-user-ID shall remain unchanged.

40309           The *seteuid()* function shall not affect the supplementary group list in any way.

40310 **RETURN VALUE**

40311           Upon successful completion, 0 shall be returned; otherwise, -1 shall be returned and *errno* set to  
40312           indicate the error.

40313 **ERRORS**

40314           The *seteuid()* function shall fail if:

40315           [EINVAL]           The value of the *uid* argument is invalid and is not supported by the  
40316           implementation.

40317           [EPERM]           The process does not have appropriate privileges and *uid* does not match the  
40318           real group ID or the saved set-group-ID.

40319 **EXAMPLES**

40320           None.

40321 **APPLICATION USAGE**

40322           None.

40323 **RATIONALE**40324           Refer to the RATIONALE section in *setuid()*.40325 **FUTURE DIRECTIONS**

40326           None.

40327 **SEE ALSO**

40328           *exec*, *getegid()*, *geteuid()*, *getgid()*, *getuid()*, *setegid()*, *setgid()*, *setregid()*, *setreuid()*, *setuid()*, the  
40329           Base Definitions volume of IEEE Std 1003.1-200x, <sys/types.h>, <unistd.h>

40330 **CHANGE HISTORY**

40331           First released in Issue 6. Derived from the IEEE P1003.1a draft standard.

40332 **NAME**

40333           setgid — set-group-ID

40334 **SYNOPSIS**

40335           #include &lt;unistd.h&gt;

40336           int setgid(gid\_t *gid*);40337 **DESCRIPTION**40338           If the process has appropriate privileges, *setgid()* shall set the real group ID, effective group ID, and the saved set-group-ID of the calling process to *gid*.40340           If the process does not have appropriate privileges, but *gid* is equal to the real group ID or the saved set-group-ID, *setgid()* shall set the effective group ID to *gid*; the real group ID and saved set-group-ID shall remain unchanged.40343           The *setgid()* function shall not affect the supplementary group list in any way.

40344           Any supplementary group IDs of the calling process shall remain unchanged.

40345 **RETURN VALUE**40346           Upon successful completion, 0 is returned. Otherwise, -1 shall be returned and *errno* set to indicate the error.40348 **ERRORS**40349           The *setgid()* function shall fail if:40350           [EINVAL]           The value of the *gid* argument is invalid and is not supported by the implementation.40352           [EPERM]           The process does not have appropriate privileges and *gid* does not match the real group ID or the saved set-group-ID.40354 **EXAMPLES**

40355           None.

40356 **APPLICATION USAGE**

40357           None.

40358 **RATIONALE**40359           Refer to the RATIONALE section in *setuid()*.40360 **FUTURE DIRECTIONS**

40361           None.

40362 **SEE ALSO**40363           *exec*, *getegid()*, *geteuid()*, *getgid()*, *getuid()*, *setegid()*, *seteuid()*, *setregid()*, *setreuid()*, *setuid()*, the Base Definitions volume of IEEE Std 1003.1-200x, <sys/types.h>, <unistd.h>40365 **CHANGE HISTORY**

40366           First released in Issue 1. Derived from Issue 1 of the SVID.

40367 **Issue 6**

40368           In the SYNOPSIS, the optional include of the &lt;sys/types.h&gt; header is removed.

40369           The following new requirements on POSIX implementations derive from alignment with the Single UNIX Specification:

- 40371           • 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.

- 40374 • Functionality associated with `_POSIX_SAVED_IDS` is now mandated. This is a FIPS  
40375 requirement.

40376 The following changes were made to align with the IEEE P1003.1a draft standard:

- 40377 • The effects of `setgid()` in processes without appropriate privileges are changed
- 40378 • A requirement that the supplementary group list is not affected is added.



40379 **NAME**

40380 setgrent — reset group database to first entry

40381 **SYNOPSIS**

40382 xSI #include &lt;grp.h&gt;

40383 void setgrent(void);

40384

40385 **DESCRIPTION**40386 Refer to *endgrent()*.

40387 **NAME**

40388           sethostent — network host database functions

40389 **SYNOPSIS**

40390           #include <netdb.h>

40391           void sethostent(int *stayopen*);

40392 **DESCRIPTION**

40393           Refer to *endhostent()*.

40394 **NAME**

40395 setitimer — set value of interval timer

40396 **SYNOPSIS**

40397 xSI #include &lt;sys/time.h&gt;

40398 int setitimer(int *which*, const struct itimerval \*restrict *value*,  
40399 struct itimerval \*restrict *ovalue*);

40400

40401 **DESCRIPTION**40402 Refer to *getitimer()*.

40403 **NAME**

40404 setjmp — set jump point for a non-local goto

40405 **SYNOPSIS**

40406 #include &lt;setjmp.h&gt;

40407 int setjmp(jmp\_buf env);

40408 **DESCRIPTION**

40409 cx The functionality described on this reference page is aligned with the ISO C standard. Any  
40410 conflict between the requirements described here and the ISO C standard is unintentional. This  
40411 volume of IEEE Std 1003.1-200x defers to the ISO C standard.

40412 A call to *setjmp()*, shall save the calling environment in its *env* argument for later use by  
40413 *longjmp()*.

40414 It is unspecified whether *setjmp()* is a macro or a function. If a macro definition is suppressed in  
40415 order to access an actual function, or a program defines an external identifier with the name  
40416 *setjmp*, the behavior is undefined.

40417 An application shall ensure that an invocation of *setjmp()* appears in one of the following  
40418 contexts only:

- 40419 • The entire controlling expression of a selection or iteration statement
- 40420 • One operand of a relational or equality operator with the other operand an integral constant  
40421 expression, with the resulting expression being the entire controlling expression of a  
40422 selection or iteration statement
- 40423 • The operand of a unary '!' operator with the resulting expression being the entire  
40424 controlling expression of a selection or iteration
- 40425 • The entire expression of an expression statement (possibly cast to **void**)

40426 If the invocation appears in any other context, the behavior is undefined.

40427 **RETURN VALUE**

40428 If the return is from a direct invocation, *setjmp()* shall return 0. If the return is from a call to  
40429 *longjmp()*, *setjmp()* shall return a non-zero value.

40430 **ERRORS**

40431 No errors are defined.

40432 **EXAMPLES**

40433 None.

40434 **APPLICATION USAGE**

40435 In general, *sigsetjmp()* is more useful in dealing with errors and interrupts encountered in a low-  
40436 level subroutine of a program.

40437 **RATIONALE**

40438 None.

40439 **FUTURE DIRECTIONS**

40440 None.

40441 **SEE ALSO**40442 *longjmp()*, *sigsetjmp()*, the Base Definitions volume of IEEE Std 1003.1-200x, <**setjmp.h**>

40443 **CHANGE HISTORY**

40444 First released in Issue 1. Derived from Issue 1 of the SVID.

40445 **Issue 6**

40446 The DESCRIPTION is updated to avoid use of the term “must” for application requirements.

40447 **NAME**40448 setkey — set encoding key (**CRYPT**)40449 **SYNOPSIS**40450 XSI `#include <stdlib.h>`40451 `void setkey(const char *key);`

40452

40453 **DESCRIPTION**

40454 The *setkey()* function provides access to an implementation-defined encoding algorithm. The  
40455 argument of *setkey()* is an array of length 64 bytes containing only the bytes with numerical  
40456 value of 0 and 1. If this string is divided into groups of 8, the low-order bit in each group is  
40457 ignored; this gives a 56-bit key which is used by the algorithm. This is the key that shall be used  
40458 with the algorithm to encode a string *block* passed to *encrypt()*.

40459 The *setkey()* function shall not change the setting of *errno* if successful. An application wishing to  
40460 check for error situations should set *errno* to 0 before calling *setkey()*. If *errno* is non-zero on  
40461 return, an error has occurred.

40462 The *setkey()* function need not be reentrant. A function that is not required to be reentrant is not  
40463 required to be thread-safe.

40464 **RETURN VALUE**

40465 No values are returned.

40466 **ERRORS**40467 The *setkey()* function shall fail if:

40468 [ENOSYS] The functionality is not supported on this implementation.

40469 **EXAMPLES**

40470 None.

40471 **APPLICATION USAGE**

40472 Decoding need not be implemented in all environments. This is related to government  
40473 restrictions in some countries on encryption and decryption routines. Historical practice has  
40474 been to ship a different version of the encryption library without the decryption feature in the  
40475 routines supplied. Thus the exported version of *encrypt()* does encoding but not decoding.

40476 **RATIONALE**

40477 None.

40478 **FUTURE DIRECTIONS**

40479 None.

40480 **SEE ALSO**40481 *crypt()*, *encrypt()*, the Base Definitions volume of IEEE Std 1003.1-200x, `<stdlib.h>`40482 **CHANGE HISTORY**

40483 First released in Issue 1. Derived from Issue 1 of the SVID.

40484 **Issue 5**40485 The DESCRIPTION is updated to indicate that *errno* is not changed if the function is successful.

## 40486 NAME

40487 setlocale — set program locale

## 40488 SYNOPSIS

40489 #include &lt;locale.h&gt;

40490 char \*setlocale(int *category*, const char \**locale*);

## 40491 DESCRIPTION

40492 CX The functionality described on this reference page is aligned with the ISO C standard. Any  
 40493 conflict between the requirements described here and the ISO C standard is unintentional. This  
 40494 volume of IEEE Std 1003.1-200x defers to the ISO C standard.

40495 The *setlocale()* function selects the appropriate piece of the program's locale, as specified by the  
 40496 *category* and *locale* arguments, and may be used to change or query the program's entire locale or  
 40497 portions thereof. The value *LC\_ALL* for *category* names the program's entire locale; other values  
 40498 for *category* name only a part of the program's locale:

40499 *LC\_COLLATE* Affects the behavior of regular expressions and the collation functions.

40500 *LC\_CTYPE* Affects the behavior of regular expressions, character classification, character  
 40501 conversion functions, and wide-character functions.

40502 CX *LC\_MESSAGES* Affects what strings are expected by commands and utilities as affirmative or  
 40503 negative responses.

40504 XSI It also affects what strings are given by commands and utilities as affirmative  
 40505 or negative responses, and the content of messages.

40506 *LC\_MONETARY* Affects the behavior of functions that handle monetary values.

40507 *LC\_NUMERIC* Affects the behavior of functions that handle numeric values.

40508 *LC\_TIME* Affects the behavior of the time conversion functions.

40509 The *locale* argument is a pointer to a character string containing the required setting of *category*.  
 40510 The contents of this string are implementation-defined. In addition, the following preset values  
 40511 of *locale* are defined for all settings of *category*:

40512 CX "POSIX" Specifies the minimal environment for C-language translation called POSIX  
 40513 locale. If *setlocale()* is not invoked, the POSIX locale is the default at entry to  
 40514 *main()*.

40515 "C" Equivalent to "POSIX".

40516 CX "" Specifies an implementation-defined native environment. This corresponds to  
 40517 the value of the associated environment variables, *LC\_\** and *LANG*; see the  
 40518 Base Definitions volume of IEEE Std 1003.1-200x, Chapter 7, Locale and the  
 40519 Base Definitions volume of IEEE Std 1003.1-200x, Chapter 8, Environment  
 40520 Variables.

40521 A null pointer Used to direct *setlocale()* to query the current internationalized environment  
 40522 and return the name of the *locale*).

40523 THR The locale state is common to all threads within a process.

## 40524 RETURN VALUE

40525 Upon successful completion, *setlocale()* shall return the string associated with the specified  
 40526 category for the new locale. Otherwise, *setlocale()* shall return a null pointer and the program's  
 40527 locale is not changed.

40528 A null pointer for *locale* causes *setlocale()* to return a pointer to the string associated with the  
 40529 *category* for the program's current locale. The program's locale shall not be changed.

40530 The string returned by *setlocale()* is such that a subsequent call with that string and its associated  
 40531 *category* shall restore that part of the program's locale. The application shall not modify the string  
 40532 returned which may be overwritten by a subsequent call to *setlocale()*.

#### 40533 ERRORS

40534 No errors are defined.

#### 40535 EXAMPLES

40536 None.

#### 40537 APPLICATION USAGE

40538 The following code illustrates how a program can initialize the international environment for  
 40539 one language, while selectively modifying the program's locale such that regular expressions  
 40540 and string operations can be applied to text recorded in a different language:

```
40541 setlocale(LC_ALL, "De");
40542 setlocale(LC_COLLATE, "Fr@dict");
```

40543 Internationalized programs must call *setlocale()* to initiate a specific language operation. This can  
 40544 be done by calling *setlocale()* as follows:

```
40545 setlocale(LC_ALL, "");
```

40546 Changing the setting of *LC\_MESSAGES* has no effect on catalogs that have already been opened  
 40547 by calls to *catopen()*.

#### 40548 RATIONALE

40549 The ISO C standard defines a collection of functions to support internationalization. One of the  
 40550 most significant aspects of these functions is a facility to set and query the *international*  
 40551 *environment*. The international environment is a repository of information that affects the  
 40552 behavior of certain functionality, namely:

- 40553 1. Character handling
- 40554 2. Collating
- 40555 3. Date/time formatting
- 40556 4. Numeric editing
- 40557 5. Monetary formatting
- 40558 6. Messaging

40559 The *setlocale()* function provides the application developer with the ability to set all or portions,  
 40560 called *categories*, of the international environment. These categories correspond to the areas of  
 40561 functionality, mentioned above. The syntax for *setlocale()* is as follows:

```
40562 char *setlocale(int category, const char *locale);
```

40563 where *category* is the name of one of following categories, namely:

```
40564     LC_COLLATE
40565     LC_CTYPE
40566     LC_MESSAGES
40567     LC_MONETARY
40568     LC_NUMERIC
40569     LC_TIME
```



40570 In addition, a special value called *LC\_ALL* directs *setlocale()* to set all categories.

40571 There are two primary uses of *setlocale()*:

- 40572 1. Querying the international environment to find out what it is set to
- 40573 2. Setting the international environment, or *locale*, to a specific value

40574 The behavior of *setlocale()* in these two areas is described below. Since it is difficult to describe  
40575 the behavior in words, examples are used to illustrate the behavior of specific uses.

40576 To query the international environment, *setlocale()* is invoked with a specific category and the  
40577 NULL pointer as the locale. The NULL pointer is a special directive to *setlocale()* that tells it to  
40578 query rather than set the international environment. The following syntax is used to query the  
40579 name of the international environment:

```
40580 setlocale({LC_ALL, LC_COLLATE, LC_CTYPE, LC_MESSAGES, LC_MONETARY, \
40581          LC_NUMERIC, LC_TIME}, (char *) NULL);
```

40582 The *setlocale()* function shall return the string corresponding to the current international  
40583 environment. This value may be used by a subsequent call to *setlocale()* to reset the international  
40584 environment to this value. However, it should be noted that the return value from *setlocale()*  
40585 may be a pointer to a static area within the function and is not guaranteed to remain unchanged  
40586 (that is, it may be modified by a subsequent call to *setlocale()*). Therefore, if the purpose of  
40587 calling *setlocale()* is to save the value of the current international environment so it can be  
40588 changed and reset later, the return value should be copied to an array of **char** in the calling  
40589 program.

40590 There are three ways to set the international environment with *setlocale()*:

40591 *setlocale(category, string)*

40592 This usage sets a specific *category* in the international environment to a specific value  
40593 corresponding to the value of the *string*. A specific example is provided below:

```
40594 setlocale(LC_ALL, "fr_FR.ISO-8859-1");
```

40595 In this example, all categories of the international environment are set to the locale  
40596 corresponding to the string "fr\_FR.ISO-8859-1", or to the French language as spoken in  
40597 France using the ISO/IEC 8859-1:1998 standard codeset.

40598 If the string does not correspond to a valid locale, *setlocale()* shall return a NULL pointer  
40599 and the international environment is not changed. Otherwise, *setlocale()* shall return the  
40600 name of the locale just set.

40601 *setlocale(category, "C")*

40602 The ISO C standard states that one locale must exist on all conforming implementations.  
40603 The name of the locale is C and corresponds to a minimal international environment needed  
40604 to support the C programming language.

40605 *setlocale(category, "")*

40606 This sets a specific category to an implementation-defined default. This corresponds to the  
40607 value of the environment variables.

#### 40608 FUTURE DIRECTIONS

40609 None.

#### 40610 SEE ALSO

40611 *exec*, *isalnum()*, *isalpha()*, *isblank()*, *iscntrl()*, *isdigit()*, *isgraph()*, *islower()*, *isprint()*, *ispunct()*,  
40612 *isspace()*, *isupper()*, *iswalnum()*, *iswalpha()*, *iswblank()*, *iswcntrl()*, *iswctype()*, *iswdigit()*,  
40613 *iswgraph()*, *iswlower()*, *iswprint()*, *iswpunct()*, *iswspace()*, *iswupper()*, *iswxdigit()*, *isxdigit()*,

40614 *localeconv()*, *mblen()*, *mbstowcs()*, *mbtowc()*, *nl\_langinfo()*, *printf()*, *scanf()*, *setlocale()*, *strcoll()*,  
40615 *strerror()*, *strfmon()*, *strtod()*, *strxfrm()*, *tolower()*, *toupper()*, *towlower()*, *towupper()*, *wscoll()*,  
40616 *wctod()*, *wcstombs()*, *wcsxfrm()*, *wctomb()*, the Base Definitions volume of IEEE Std 1003.1-200x,  
40617 **<langinfo.h>**, **<locale.h>**

40618 **CHANGE HISTORY**

40619 First released in Issue 3.

40620 **Issue 5**

40621 The DESCRIPTION is updated for alignment with the POSIX Threads Extension.

40622 **Issue 6**

40623 Extensions beyond the ISO C standard are now marked.

40624 The DESCRIPTION is updated to avoid use of the term “must” for application requirements.

40625 **NAME**

40626           setlogmask — set log priority mask

40627 **SYNOPSIS**

40628 XSI       #include &lt;syslog.h&gt;

40629           int setlogmask(int maskpri);

40630

40631 **DESCRIPTION**40632           Refer to *closelog()*.

# setnetent()

40633 **NAME**

40634           setnetent — network database function

40635 **SYNOPSIS**

40636           #include <netdb.h>

40637           void setnetent(int stayopen);

40638 **DESCRIPTION**

40639           Refer to *endnetent()*.

40640 **NAME**

40641 setpgid — set process group ID for job control

40642 **SYNOPSIS**

40643 #include &lt;unistd.h&gt;

40644 int setpgid(pid\_t pid, pid\_t pgid);

40645 **DESCRIPTION**

40646 The *setpgid()* function shall either join an existing process group or create a new process group  
 40647 within the session of the calling process. The process group ID of a session leader shall not  
 40648 change. Upon successful completion, the process group ID of the process with a process ID that  
 40649 matches *pid* shall be set to *pgid*. As a special case, if *pid* is 0, the process ID of the calling process  
 40650 shall be used. Also, if *pgid* is 0, the process group ID of the indicated process shall be used.

40651 **RETURN VALUE**

40652 Upon successful completion, *setpgid()* shall return 0; otherwise, -1 shall be returned and *errno*  
 40653 shall be set to indicate the error.

40654 **ERRORS**40655 The *setpgid()* function shall fail if:

40656 [EACCES] The value of the *pid* argument matches the process ID of a child process of the  
 40657 calling process and the child process has successfully executed one of the *exec*  
 40658 functions.

40659 [EINVAL] The value of the *pgid* argument is less than 0, or is not a value supported by  
 40660 the implementation.

40661 [EPERM] The process indicated by the *pid* argument is a session leader.

40662 [EPERM] The value of the *pid* argument matches the process ID of a child process of the  
 40663 calling process and the child process is not in the same session as the calling  
 40664 process.

40665 [EPERM] The value of the *pgid* argument is valid but does not match the process ID of  
 40666 the process indicated by the *pid* argument and there is no process with a  
 40667 process group ID that matches the value of the *pgid* argument in the same  
 40668 session as the calling process.

40669 [ESRCH] The value of the *pid* argument does not match the process ID of the calling  
 40670 process or of a child process of the calling process.

40671 **EXAMPLES**

40672 None.

40673 **APPLICATION USAGE**

40674 None.

40675 **RATIONALE**

40676 The *setpgid()* function shall group processes together for the purpose of signaling, placement in  
 40677 foreground or background, and other job control actions.

40678 The *setpgid()* function is similar to the *setpgrp()* function of 4.2 BSD, except that 4.2 BSD allowed  
 40679 the specified new process group to assume any value. This presents certain security problems  
 40680 and is more flexible than necessary to support job control.

40681 To provide tighter security, *setpgid()* only allows the calling process to join a process group  
 40682 already in use inside its session or create a new process group whose process group ID was  
 40683 equal to its process ID.

40684 When a job control shell spawns a new job, the processes in the job must be placed into a new  
40685 process group via *setpgid()*. There are two timing constraints involved in this action:

- 40686 1. The new process must be placed in the new process group before the appropriate program  
40687 is launched via one of the *exec* functions.
- 40688 2. The new process must be placed in the new process group before the shell can correctly  
40689 send signals to the new process group.

40690 To address these constraints, the following actions are performed. The new processes call  
40691 *setpgid()* to alter their own process groups after *fork()* but before *exec*. This satisfies the first  
40692 constraint. Under 4.3 BSD, the second constraint is satisfied by the synchronization property of  
40693 *vfork()*; that is, the shell is suspended until the child has completed the *exec*, thus ensuring that  
40694 the child has completed the *setpgid()*. A new version of *fork()* with this same synchronization  
40695 property was considered, but it was decided instead to merely allow the parent shell process to  
40696 adjust the process group of its child processes via *setpgid()*. Both timing constraints are now  
40697 satisfied by having both the parent shell and the child attempt to adjust the process group of the  
40698 child process; it does not matter which succeeds first.

40699 Since it would be confusing to an application to have its process group change after it began  
40700 executing (that is, after *exec*), and because the child process would already have adjusted its  
40701 process group before this, the [EACCES] error was added to disallow this.

40702 One non-obvious use of *setpgid()* is to allow a job control shell to return itself to its original  
40703 process group (the one in effect when the job control shell was executed). A job control shell  
40704 does this before returning control back to its parent when it is terminating or suspending itself as  
40705 a way of restoring its job control “state” back to what its parent would expect. (Note that the  
40706 original process group of the job control shell typically matches the process group of its parent,  
40707 but this is not necessarily always the case.)

#### 40708 FUTURE DIRECTIONS

40709 None.

#### 40710 SEE ALSO

40711 *exec*, *getpgrp()*, *setsid()*, *tcsetpgrp()*, the Base Definitions volume of IEEE Std 1003.1-200x,  
40712 `<sys/types.h>`, `<unistd.h>`

#### 40713 CHANGE HISTORY

40714 First released in Issue 3.

40715 Entry included for alignment with the POSIX.1-1988 standard.

#### 40716 Issue 6

40717 In the SYNOPSIS, the optional include of the `<sys/types.h>` header is removed.

40718 The following new requirements on POSIX implementations derive from alignment with the  
40719 Single UNIX Specification:

- 40720 • The requirement to include `<sys/types.h>` has been removed. Although `<sys/types.h>` was  
40721 required for conforming implementations of previous POSIX specifications, it was not  
40722 required for UNIX applications.
- 40723 • The *setpgid()* function is mandatory since `_POSIX_JOB_CONTROL` is required to be defined  
40724 in this issue. This is a FIPS requirement.

40725 **NAME**

40726           setpgrp — set process group ID

40727 **SYNOPSIS**

40728 XSI       #include &lt;unistd.h&gt;

40729           pid\_t setpgrp(void);

40730

40731 **DESCRIPTION**

40732           If the calling process is not already a session leader, *setpgrp()* sets the process group ID of the  
40733           calling process to the process ID of the calling process. If *setpgrp()* creates a new session, then  
40734           the new session has no controlling terminal.

40735           The *setpgrp()* function has no effect when the calling process is a session leader.

40736 **RETURN VALUE**40737           Upon completion, *setpgrp()* shall return the process group ID.40738 **ERRORS**

40739           No errors are defined.

40740 **EXAMPLES**

40741           None.

40742 **APPLICATION USAGE**

40743           None.

40744 **RATIONALE**

40745           None.

40746 **FUTURE DIRECTIONS**

40747           None.

40748 **SEE ALSO**

40749           *exec*, *fork()*, *getpid()*, *getsid()*, *kill()*, *setpgid()*, *setsid()*, the Base Definitions volume of  
40750           IEEE Std 1003.1-200x, <unistd.h>

40751 **CHANGE HISTORY**

40752           First released in Issue 4, Version 2.

40753 **Issue 5**

40754           Moved from X/OPEN UNIX extension to BASE.

40755 **NAME**

40756           setpriority — set the nice value

40757 **SYNOPSIS**

40758 XSI       #include <sys/resource.h>

40759           int setpriority(int *which*, id\_t *who*, int *nice*);

40760

40761 **DESCRIPTION**

40762           Refer to *getpriority()*.



40763 **NAME**

40764           setprotoent — network protocol database functions

40765 **SYNOPSIS**

40766           #include <netdb.h>

40767           void setprotoent(int *stayopen*);

40768 **DESCRIPTION**

40769           Refer to *endprotoent()*.

40770 **NAME**

40771 setpwent — user database function

40772 **SYNOPSIS**

40773 XSI #include <pwd.h>

40774 void setpwent(void);

40775

40776 **DESCRIPTION**

40777 Refer to *endpwent()*.

40778 **NAME**

40779 setregid — set real and effective group IDs

40780 **SYNOPSIS**

40781 XSI #include &lt;unistd.h&gt;

40782 int setregid(gid\_t rgid, gid\_t egid);

40783

40784 **DESCRIPTION**40785 The *setregid()* function shall set the real and effective group IDs of the calling process. |40786 If *rgid* is  $-1$ , the real group ID shall not be changed; if *egid* is  $-1$ , the effective group ID shall not  
40787 be changed.

40788 The real and effective group IDs may be set to different values in the same call.

40789 Only a process with appropriate privileges can set the real group ID and the effective group ID  
40790 to any valid value.40791 A non-privileged process can set either the real group ID to the saved set-group-ID from one of  
40792 the *exec* family of functions, or the effective group ID to the saved set-group-ID or the real group  
40793 ID.

40794 Any supplementary group IDs of the calling process remain unchanged.

40795 **RETURN VALUE**40796 Upon successful completion, 0 shall be returned. Otherwise,  $-1$  shall be returned and *errno* set to  
40797 indicate the error, and neither of the group IDs are changed.40798 **ERRORS**40799 The *setregid()* function shall fail if:40800 [EINVAL] The value of the *rgid* or *egid* argument is invalid or out-of-range.40801 [EPERM] The process does not have appropriate privileges and a change other than  
40802 changing the real group ID to the saved set-group-ID, or changing the  
40803 effective group ID to the real group ID or the saved set-group-ID, was  
40804 requested.40805 **EXAMPLES**

40806 None.

40807 **APPLICATION USAGE**40808 If a set-group-ID process sets its effective group ID to its real group ID, it can still set its effective  
40809 group ID back to the saved set-group-ID.40810 **RATIONALE**

40811 None.

40812 **FUTURE DIRECTIONS**

40813 None.

40814 **SEE ALSO**40815 *exec*, *getegid()*, *geteuid()*, *getgid()*, *getuid()*, *setegid()*, *seteuid()*, *setgid()*, *setreuid()*, *setuid()*, the  
40816 Base Definitions volume of IEEE Std 1003.1-200x, <unistd.h>40817 **CHANGE HISTORY**

40818 First released in Issue 4, Version 2.

40819 **Issue 5**

40820 Moved from X/OPEN UNIX extension to BASE.

40821 The DESCRIPTION is updated to indicate that the saved set-group-ID can be set by any of the  
40822 *exec* family of functions, not just *execev()*.

40823 **NAME**

40824 setreuid — set real and effective user IDs

40825 **SYNOPSIS**40826 XSI `#include <unistd.h>`40827 `int setreuid(uid_t ruid, uid_t euid);`

40828

40829 **DESCRIPTION**

40830 The `setreuid()` function shall set the real and effective user IDs of the current process to the  
 40831 values specified by the `ruid` and `euid` arguments. If `ruid` or `euid` is `-1`, the corresponding effective  
 40832 or real user ID of the current process shall be left unchanged.

40833 A process with appropriate privileges can set either ID to any value. An unprivileged process  
 40834 can only set the effective user ID if the `euid` argument is equal to either the real, effective, or  
 40835 saved user ID of the process.

40836 It is unspecified whether a process without appropriate privileges is permitted to change the real  
 40837 user ID to match the current real, effective, or saved set-user-ID of the process.

40838 **RETURN VALUE**

40839 Upon successful completion, 0 shall be returned. Otherwise, `-1` shall be returned and `errno` set to  
 40840 indicate the error.

40841 **ERRORS**40842 The `setreuid()` function shall fail if:40843 [EINVAL] The value of the `ruid` or `euid` argument is invalid or out-of-range.

40844 [EPERM] The current process does not have appropriate privileges, and either an  
 40845 attempt was made to change the effective user ID to a value other than the  
 40846 real user ID or the saved set-user-ID or an attempt was made to change the  
 40847 real user ID to a value not permitted by the implementation.

40848 **EXAMPLES**40849 **Setting the Effective User ID to the Real User ID**

40850 The following example sets the effective user ID of the calling process to the real user ID, so that  
 40851 files created later will be owned by the current user.

```
40852 #include <unistd.h>
40853 #include <sys/types.h>
40854 ...
40855 setreuid(getuid(), getuid());
40856 ...
```

40857 **APPLICATION USAGE**

40858 None.

40859 **RATIONALE**

40860 None.

40861 **FUTURE DIRECTIONS**

40862 None.

40863 **SEE ALSO**

40864 *getegid()*, *geteuid()*, *getgid()*, *getuid()*, *setegid()*, *seteuid()*, *setgid()*, *setregid()*, *setuid()*, the Base  
40865 Definitions volume of IEEE Std 1003.1-200x, <**unistd.h**>

40866 **CHANGE HISTORY**

40867 First released in Issue 4, Version 2.

40868 **Issue 5**

40869 Moved from X/OPEN UNIX extension to BASE.

40870 **NAME**

40871 setrlimit — control maximum resource consumption

40872 **SYNOPSIS**

40873 XSI #include &lt;sys/resource.h&gt;

40874 int setrlimit(int resource, const struct rlimit \*rlp);

40875

40876 **DESCRIPTION**40877 Refer to *getrlimit()*.

40878 **NAME**

40879       setservent — network services database functions

40880 **SYNOPSIS**

40881       #include <netdb.h>

40882       void setservent(int *stayopen*);

40883 **DESCRIPTION**

40884       Refer to *endservent()*.



40885 **NAME**

40886 setsid — create session and set process group ID

40887 **SYNOPSIS**

40888 #include &lt;unistd.h&gt;

40889 pid\_t setsid(void);

40890 **DESCRIPTION**

40891 The *setsid()* function shall create a new session, if the calling process is not a process group  
40892 leader. Upon return the calling process shall be the session leader of this new session, shall be  
40893 the process group leader of a new process group, and shall have no controlling terminal. The  
40894 process group ID of the calling process shall be set equal to the process ID of the calling process.  
40895 The calling process shall be the only process in the new process group and the only process in  
40896 the new session.

40897 **RETURN VALUE**

40898 Upon successful completion, *setsid()* shall return the value of the new process group ID of the  
40899 calling process. Otherwise, it shall return (**pid\_t**)-1 and set *errno* to indicate the error.

40900 **ERRORS**40901 The *setsid()* function shall fail if:

40902 [EPERM] The calling process is already a process group leader, or the process group ID  
40903 of a process other than the calling process matches the process ID of the  
40904 calling process.

40905 **EXAMPLES**

40906 None.

40907 **APPLICATION USAGE**

40908 None.

40909 **RATIONALE**

40910 The *setsid()* function is similar to the *setpgrp()* function of System V. System V, without job  
40911 control, groups processes into process groups and creates new process groups via *setpgrp()*; only  
40912 one process group may be part of a login session.

40913 Job control allows multiple process groups within a login session. In order to limit job control  
40914 actions so that they can only affect processes in the same login session, this volume of  
40915 IEEE Std 1003.1-200x adds the concept of a session that is created via *setsid()*. The *setsid()*  
40916 function also creates the initial process group contained in the session. Additional process  
40917 groups can be created via the *setpgid()* function. A System V process group would correspond to  
40918 a POSIX System Interfaces session containing a single POSIX process group. Note that this  
40919 function requires that the calling process not be a process group leader. The usual way to ensure  
40920 this is true is to create a new process with *fork()* and have it call *setsid()*. The *fork()* function  
40921 guarantees that the process ID of the new process does not match any existing process group ID.

40922 **FUTURE DIRECTIONS**

40923 None.

40924 **SEE ALSO**

40925 *getsid()*, *setpgid()*, *setpgrp()*, the Base Definitions volume of IEEE Std 1003.1-200x, <sys/types.h>,  
40926 <unistd.h>

40927 **CHANGE HISTORY**

40928 First released in Issue 3.

40929 Entry included for alignment with the POSIX.1-1988 standard.

40930 **Issue 6**40931 In the SYNOPSIS, the optional include of the `<sys/types.h>` header is removed.40932 The following new requirements on POSIX implementations derive from alignment with the  
40933 Single UNIX Specification:

- 40934
- The requirement to include `<sys/types.h>` has been removed. Although `<sys/types.h>` was  
40935 required for conforming implementations of previous POSIX specifications, it was not  
40936 required for UNIX applications.

## 40937 NAME

40938 setsockopt — set the socket options

## 40939 SYNOPSIS

40940 #include &lt;sys/socket.h&gt;

```
40941 int setsockopt(int socket, int level, int option_name,
40942               const void *option_value, socklen_t option_len);
```

## 40943 DESCRIPTION

40944 The *setsockopt()* function shall set the option specified by the *option\_name* argument, at the |  
 40945 protocol level specified by the *level* argument, to the value pointed to by the *option\_value*  
 40946 argument for the socket associated with the file descriptor specified by the *socket* argument.

40947 The *level* argument specifies the protocol level at which the option resides. To set options at the  
 40948 socket level, specify the *level* argument as SOL\_SOCKET. To set options at other levels, supply  
 40949 the appropriate *level* identifier for the protocol controlling the option. For example, to indicate  
 40950 that an option is interpreted by the TCP (Transport Control Protocol), set *level* to IPPROTO\_TCP  
 40951 as defined in the <netinet/in.h> header.

40952 The *option\_name* argument specifies a single option to set. The *option\_name* argument and any  
 40953 specified options are passed uninterpreted to the appropriate protocol module for  
 40954 interpretations. The <sys/socket.h> header defines the socket-level options. The options are as  
 40955 follows:

40956 SO\_DEBUG Turns on recording of debugging information. This option enables or  
 40957 disables debugging in the underlying protocol modules. This option takes  
 40958 an **int** value. This is a Boolean option.

40959 SO\_BROADCAST Permits sending of broadcast messages, if this is supported by the  
 40960 protocol. This option takes an **int** value. This is a Boolean option.

40961 SO\_REUSEADDR Specifies that the rules used in validating addresses supplied to *bind()*  
 40962 should allow reuse of local addresses, if this is supported by the protocol.  
 40963 This option takes an **int** value. This is a Boolean option.

40964 SO\_KEEPALIVE Keeps connections active by enabling the periodic transmission of  
 40965 messages, if this is supported by the protocol. This option takes an **int**  
 40966 value.

40967 If the connected socket fails to respond to these messages, the connection  
 40968 is broken and threads writing to that socket are notified with a SIGPIPE  
 40969 signal.

40970 This is a Boolean option.

40971 SO\_LINGER Lingers on a *close()* if data is present. This option controls the action  
 40972 taken when unsent messages queue on a socket and *close()* is performed. |  
 40973 If SO\_LINGER is set, the system shall block the process during *close()* |  
 40974 until it can transmit the data or until the time expires. If SO\_LINGER is  
 40975 not specified, and *close()* is issued, the system handles the call in a way  
 40976 that allows the process to continue as quickly as possible. This option  
 40977 takes a **linger** structure, as defined in the <sys/socket.h> header, to  
 40978 specify the state of the option and linger interval.

40979 SO\_OOBINLINE Leaves received out-of-band data (data marked urgent) inline. This  
 40980 option takes an **int** value. This is a Boolean option.

40981	SO_SNDBUF	Sets send buffer size. This option takes an <b>int</b> value.
40982	SO_RCVBUF	Sets receive buffer size. This option takes an <b>int</b> value.
40983	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.
40984		
40985		
40986		
40987		
40988	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 from 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.
40989		
40990		
40991		
40992		
40993		
40994		
40995		
40996	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.
40997		
40998		
40999		
41000		
41001		
41002		
41003		
41004		
41005		
41006	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.
41007		
41008		
41009		
41010		
41011	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.
41012		
41013		
41014		
41015		
41016		
41017		
41018		
41019		
41020		
41020		Options at other protocol levels vary in format and name.
41021		
41021	<b>RETURN VALUE</b>	
41022		Upon successful completion, <i>setsockopt()</i> shall return 0. Otherwise, -1 shall be returned and <i>errno</i> set to indicate the error.
41023		
41024	<b>ERRORS</b>	
41025		The <i>setsockopt()</i> function shall fail if:
41026	[EBADF]	The <i>socket</i> argument is not a valid file descriptor.

- 41027 [EDOM] The send and receive timeout values are too big to fit into the timeout fields in  
41028 the socket structure.
- 41029 [EINVAL] The specified option is invalid at the specified socket level or the socket has  
41030 been shut down.
- 41031 [EISCONN] The socket is already connected, and a specified option cannot be set while the  
41032 socket is connected.
- 41033 [ENOPROTOOPT]  
41034 The option is not supported by the protocol.
- 41035 [ENOTSOCK] The *socket* argument does not refer to a socket.
- 41036 The *setsockopt()* function may fail if:
- 41037 [ENOMEM] There was insufficient memory available for the operation to complete.
- 41038 [ENOBUFS] Insufficient resources are available in the system to complete the call.
- 41039 **EXAMPLES**
- 41040 None.
- 41041 **APPLICATION USAGE**
- 41042 The *setsockopt()* function provides an application program with the means to control socket  
41043 behavior. An application program can use *setsockopt()* to allocate buffer space, control timeouts,  
41044 or permit socket data broadcasts. The `<sys/socket.h>` header defines the socket-level options  
41045 available to *setsockopt()*.
- 41046 Options may exist at multiple protocol levels. The `SO_` options are always present at the  
41047 uppermost socket level.
- 41048 **RATIONALE**
- 41049 None.
- 41050 **FUTURE DIRECTIONS**
- 41051 None.
- 41052 **SEE ALSO**
- 41053 Section 2.10 (on page 508), *bind()*, *endprotoent()*, *getsockopt()*, *socket()*, the Base Definitions  
41054 volume of IEEE Std 1003.1-200x, `<netinet/in.h>`, `<sys/socket.h>`
- 41055 **CHANGE HISTORY**
- 41056 First released in Issue 6. Derived from the XNS, Issue 5.2 specification.

41057 **NAME**

41058            setstate — switch pseudo-random number generator state arrays

41059 **SYNOPSIS**

41060 xSI        #include <stdlib.h>

41061            char \*setstate(const char \*state);

41062

41063 **DESCRIPTION**

41064            Refer to *initstate()*.

41065 **NAME**

41066           setuid — set user ID

41067 **SYNOPSIS**

41068           #include &lt;unistd.h&gt;

41069           int setuid(uid\_t uid);

41070 **DESCRIPTION**41071           If the process has appropriate privileges, *setuid()* shall set the real user ID, effective user ID, and  
41072           the saved set-user-ID of the calling process to *uid*.41073           If the process does not have appropriate privileges, but *uid* is equal to the real user ID or the  
41074           saved set-user-ID, *setuid()* shall set the effective user ID to *uid*; the real user ID and saved set-  
41075           user-ID shall remain unchanged.41076           The *setuid()* function shall not affect the supplementary group list in any way.41077 **RETURN VALUE**41078           Upon successful completion, 0 shall be returned. Otherwise, -1 shall be returned and *errno* set to  
41079           indicate the error.41080 **ERRORS**41081           The *setuid()* function shall fail, return -1, and set *errno* to the corresponding value if one or more  
41082           of the following are true:41083           [EINVAL]           The value of the *uid* argument is invalid and not supported by the  
41084           implementation.41085           [EPERM]           The process does not have appropriate privileges and *uid* does not match the  
41086           real user ID or the saved set-user-ID.41087 **EXAMPLES**

41088           None.

41089 **APPLICATION USAGE**

41090           None.

41091 **RATIONALE**41092           The various behaviors of the *setuid()* and *setgid()* functions when called by non-privileged  
41093           processes reflect the behavior of different historical implementations. For portability, it is  
41094           recommended that new non-privileged applications use the *seteuid()* and *setegid()* functions  
41095           instead.41096           The saved set-user-ID capability allows a program to regain the effective user ID established at  
41097           the last *exec* call. Similarly, the saved set-group-ID capability allows a program to regain the  
41098           effective group ID established at the last *exec* call. These capabilities are derived from System V.  
41099           Without them, a program might have to run as superuser in order to perform the same |  
41100           functions, because superuser can write on the user's files. This is a problem because such a |  
41101           program can write on any user's files, and so must be carefully written to emulate the |  
41102           permissions of the calling process properly. In System V, these capabilities have traditionally |  
41103           been implemented only via the *setuid()* and *setgid()* functions for non-privileged processes. The  
41104           fact that the behavior of those functions was different for privileged processes made them  
41105           difficult to use. The POSIX.1-1990 standard defined the *setuid()* function to behave differently  
41106           for privileged and unprivileged users. When the caller had the appropriate privilege, the  
41107           function set the calling process' real user ID, effective user ID, and saved set-user ID on |  
41108           implementations that supported it. When the caller did not have the appropriate privilege, the  
41109           function set only the effective user ID, subject to permission checks. The former use is generally  
41110           needed for utilities like *login* and *su*, which are not conforming applications and thus outside the |

41111 scope of IEEE Std 1003.1-200x. These utilities wish to change the user ID irrevocably to a new |  
41112 value, generally that of an unprivileged user. The latter use is needed for conforming |  
41113 applications that are installed with the set-user-ID bit and need to perform operations using the |  
41114 real user ID.

41115 IEEE Std 1003.1-200x augments the latter functionality with a mandatory feature named |  
41116 `_POSIX_SAVED_IDS`. This feature permits a set-user-ID application to switch its effective user |  
41117 ID back and forth between the values of its *exec*-time real user ID and effective user ID. |  
41118 Unfortunately, the POSIX.1-1990 standard did not permit a conforming application using this |  
41119 feature to work properly when it happened to be executed with the (implementation-defined) |  
41120 appropriate privilege. Furthermore, the application did not even have a means to tell whether it |  
41121 had this privilege. Since the saved set-user-ID feature is quite desirable for applications, as |  
41122 evidenced by the fact that NIST required it in FIPS 151-2, it has been mandated by |  
41123 IEEE Std 1003.1-200x. However, there are implementors who have been reluctant to support it |  
41124 given the limitation described above.

41125 The 4.3BSD system handles the problem by supporting separate functions: *setuid()* (which |  
41126 always sets both the real and effective user IDs, like *setuid()* in IEEE Std 1003.1-200x for |  
41127 privileged users), and *seteuid()* (which always sets just the effective user ID, like *setuid()* in |  
41128 IEEE Std 1003.1-200x for non-privileged users). This separation of functionality into distinct |  
41129 functions seems desirable. 4.3BSD does not support the saved set-user-ID feature. It supports |  
41130 similar functionality of switching the effective user ID back and forth via *setreuid()*, which |  
41131 permits reversing the real and effective user IDs. This model seems less desirable than the saved |  
41132 set-user-ID because the real user ID changes as a side effect. The current 4.4BSD includes saved |  
41133 effective IDs and uses them for *seteuid()* and *setegid()* as described above. The *setreuid()* and |  
41134 *setregid()* functions will be deprecated or removed.

41135 The solution here is:

- 41136 • Require that all implementations support the functionality of the saved set-user-ID, which is |  
41137 set by the *exec* functions and by privileged calls to *setuid()*.
- 41138 • Add the *seteuid()* and *setegid()* functions as portable alternatives to *setuid()* and *setgid()* for |  
41139 non-privileged and privileged processes.

41140 Historical systems have provided two mechanisms for a set-user-ID process to change its |  
41141 effective user ID to be the same as its real user ID in such a way that it could return to the |  
41142 original effective user ID: the use of the *setuid()* function in the presence of a saved set-user-ID, |  
41143 or the use of the BSD *setreuid()* function, which was able to swap the real and effective user IDs. |  
41144 The changes included in IEEE Std 1003.1-200x provide a new mechanism using *seteuid()* |  
41145 in conjunction with a saved set-user-ID. Thus, all implementations with the new *seteuid()* |  
41146 mechanism will have a saved set-user-ID for each process, and most of the behavior controlled |  
41147 by `_POSIX_SAVED_IDS` has been changed to agree with the case where the option was defined. |  
41148 The *kill()* function is an exception. Implementors of the new *seteuid()* mechanism will generally |  
41149 be required to maintain compatibility with the older mechanisms previously supported by their |  
41150 systems. However, compatibility with this use of *setreuid()* and with the `_POSIX_SAVED_IDS` |  
41151 behavior of *kill()* is unfortunately complicated. If an implementation with a saved set-user-ID |  
41152 allows a process to use *setreuid()* to swap its real and effective user IDs, but were to leave the |  
41153 saved set-user-ID unmodified, the process would then have an effective user ID equal to the |  
41154 original real user ID, and both real and saved set-user-ID would be equal to the original effective |  
41155 user ID. In that state, the real user would be unable to kill the process, even though the effective |  
41156 user ID of the process matches that of the real user, if the *kill()* behavior of `_POSIX_SAVED_IDS` |  
41157 was used. This is obviously not acceptable. The alternative choice, which is used in at least one |  
41158 implementation, is to change the saved set-user-ID to the effective user ID during most calls to |  
41159 *setreuid()*. The standard developers considered that alternative to be less correct than the



41160 retention of the old behavior of *kill()* in such systems. Current conforming applications shall  
41161 accommodate either behavior from *kill()*, and there appears to be no strong reason for *kill()* to  
41162 check the saved set-user-ID rather than the effective user ID.

#### 41163 FUTURE DIRECTIONS

41164 None.

#### 41165 SEE ALSO

41166 *exec*, *getegid()*, *geteuid()*, *getgid()*, *getuid()*, *setegid()*, *seteuid()*, *setgid()*, *setregid()*, *setreuid()*, the  
41167 Base Definitions volume of IEEE Std 1003.1-200x, `<sys/types.h>`, `<unistd.h>`

#### 41168 CHANGE HISTORY

41169 First released in Issue 1. Derived from Issue 1 of the SVID.

#### 41170 Issue 6

41171 In the SYNOPSIS, the optional include of the `<sys/types.h>` header is removed.

41172 The following new requirements on POSIX implementations derive from alignment with the  
41173 Single UNIX Specification:

- 41174 • The requirement to include `<sys/types.h>` has been removed. Although `<sys/types.h>` was  
41175 required for conforming implementations of previous POSIX specifications, it was not  
41176 required for UNIX applications.
- 41177 • The functionality associated with `_POSIX_SAVED_IDS` is now mandatory. This is a FIPS  
41178 requirement.

41179 The following changes were made to align with the IEEE P1003.1a draft standard:

- 41180 • The effects of *setuid()* in processes without appropriate privileges are changed.
- 41181 • A requirement that the supplementary group list is not affected is added.

41182 **NAME**

41183           setutxent — reset user accounting database to first entry

41184 **SYNOPSIS**

41185 XSI       #include <utmpx.h>

41186           void setutxent(void);

41187

41188 **DESCRIPTION**

41189           Refer to *endutxent()*.

41190 **NAME**

41191 setvbuf — assign buffering to a stream

41192 **SYNOPSIS**

41193 #include &lt;stdio.h&gt;

41194 int setvbuf(FILE \*restrict stream, char \*restrict buf, int type,  
41195 size\_t size);41196 **DESCRIPTION**41197 **CX** The functionality described on this reference page is aligned with the ISO C standard. Any  
41198 conflict between the requirements described here and the ISO C standard is unintentional. This  
41199 volume of IEEE Std 1003.1-200x defers to the ISO C standard.41200 The *setvbuf()* function may be used after the stream pointed to by *stream* is associated with an  
41201 open file but before any other operation (other than an unsuccessful call to *setvbuf()*) is  
41202 performed on the stream. The argument *type* determines how *stream* shall be buffered, as  
41203 follows:

- 41204 • {\_IOFBF} shall cause input/output to be fully buffered.
- 41205 • {\_IOLBF} shall cause input/output to be line buffered.
- 41206 • {\_IONBF} shall cause input/output to be unbuffered.

41207 If *buf* is not a null pointer, the array it points to may be used instead of a buffer allocated by  
41208 *setvbuf()* and the argument *size* specifies the size of the array; otherwise, *size* may determine the  
41209 size of a buffer allocated by the *setvbuf()* function. The contents of the array at any time are  
41210 unspecified. |

41211 For information about streams, see Section 2.5 (on page 484).

41212 **RETURN VALUE**41213 Upon successful completion, *setvbuf()* shall return 0. Otherwise, it shall return a non-zero value  
41214 **CX** if an invalid value is given for *type* or if the request cannot be honored, and may set *errno* to  
41215 indicate the error.41216 **ERRORS**41217 The *setvbuf()* function may fail if:41218 **CX** [EBADF] The file descriptor underlying *stream* is not valid.41219 **EXAMPLES**

41220 None.

41221 **APPLICATION USAGE**41222 A common source of error is allocating buffer space as an “automatic” variable in a code block,  
41223 and then failing to close the stream in the same block.41224 With *setvbuf()*, allocating a buffer of *size* bytes does not necessarily imply that all of *size* bytes are  
41225 used for the buffer area.41226 Applications should note that many implementations only provide line buffering on input from  
41227 terminal devices.41228 **RATIONALE**

41229 None.

41230 **FUTURE DIRECTIONS**

41231 None.

41232 **SEE ALSO**

41233 *fopen()*, *setbuf()*, the Base Definitions volume of IEEE Std 1003.1-200x, <**stdio.h**>

41234 **CHANGE HISTORY**

41235 First released in Issue 1. Derived from Issue 1 of the SVID.

41236 **Issue 6**

41237 Extensions beyond the ISO C standard are now marked.

41238 The *setvbuf()* prototype is updated for alignment with the ISO/IEC 9899: 1999 standard.

## 41239 NAME

41240 shm\_open — open a shared memory object (**REALTIME**)

## 41241 SYNOPSIS

41242 SHM #include &lt;sys/mman.h&gt;

41243 int shm\_open(const char \*name, int oflag, mode\_t mode);

41244

## 41245 DESCRIPTION

41246 The *shm\_open()* function shall establish a connection between a shared memory object and a file  
 41247 descriptor. It shall create an open file description that refers to the shared memory object and a  
 41248 file descriptor that refers to that open file description. The file descriptor is used by other  
 41249 functions to refer to that shared memory object. The *name* argument points to a string naming a  
 41250 shared memory object. It is unspecified whether the name appears in the file system and is  
 41251 visible to other functions that take pathnames as arguments. The *name* argument conforms to the  
 41252 construction rules for a pathname. If *name* begins with the slash character, then processes calling  
 41253 *shm\_open()* with the same value of *name* refer to the same shared memory object, as long as that  
 41254 name has not been removed. If *name* does not begin with the slash character, the effect is  
 41255 implementation-defined. The interpretation of slash characters other than the leading slash  
 41256 character in *name* is implementation-defined.

41257 If successful, *shm\_open()* shall return a file descriptor for the shared memory object that is the  
 41258 lowest numbered file descriptor not currently open for that process. The open file description is  
 41259 new, and therefore the file descriptor does not share it with any other processes. It is unspecified  
 41260 whether the file offset is set. The FD\_CLOEXEC file descriptor flag associated with the new file  
 41261 descriptor is set.

41262 The file status flags and file access modes of the open file description are according to the value  
 41263 of *oflag*. The *oflag* argument is the bitwise-inclusive OR of the following flags defined in the  
 41264 <fcntl.h> header. Applications specify exactly one of the first two values (access modes) below  
 41265 in the value of *oflag*:

41266 O\_RDONLY Open for read access only.

41267 O\_RDWR Open for read or write access.

41268 Any combination of the remaining flags may be specified in the value of *oflag*:

41269 O\_CREAT If the shared memory object exists, this flag has no effect, except as noted  
 41270 under O\_EXCL below. Otherwise, the shared memory object is created; the user ID of the shared memory object shall be set to the effective user ID of the  
 41271 process; the group ID of the shared memory object is set to a system default group ID or to the effective group ID of the process. The permission bits of the  
 41272 shared memory object shall be set to the value of the *mode* argument except  
 41273 those set in the file mode creation mask of the process. When bits in *mode*  
 41274 other than the file permission bits are set, the effect is unspecified. The *mode*  
 41275 argument does not affect whether the shared memory object is opened for  
 41276 reading, for writing, or for both. The shared memory object has a size of zero.

41279 O\_EXCL If O\_EXCL and O\_CREAT are set, *shm\_open()* fails if the shared memory  
 41280 object exists. The check for the existence of the shared memory object and the  
 41281 creation of the object if it does not exist is atomic with respect to other  
 41282 processes executing *shm\_open()* naming the same shared memory object with  
 41283 O\_EXCL and O\_CREAT set. If O\_EXCL is set and O\_CREAT is not set, the  
 41284 result is undefined.

41285 O\_TRUNC If the shared memory object exists, and it is successfully opened O\_RDWR,  
 41286 the object shall be truncated to zero length and the mode and owner shall be  
 41287 unchanged by this function call. The result of using O\_TRUNC with  
 41288 O\_RDONLY is undefined.

41289 When a shared memory object is created, the state of the shared memory object, including all  
 41290 data associated with the shared memory object, persists until the shared memory object is  
 41291 unlinked and all other references are gone. It is unspecified whether the name and shared  
 41292 memory object state remain valid after a system reboot.

#### 41293 RETURN VALUE

41294 Upon successful completion, the *shm\_open()* function shall return a non-negative integer  
 41295 representing the lowest numbered unused file descriptor. Otherwise, it shall return  $-1$  and set  
 41296 *errno* to indicate the error.

#### 41297 ERRORS

41298 The *shm\_open()* function shall fail if:

41299 [EACCES] The shared memory object exists and the permissions specified by *oflag* are  
 41300 denied, or the shared memory object does not exist and permission to create  
 41301 the shared memory object is denied, or O\_TRUNC is specified and write  
 41302 permission is denied.

41303 [EEXIST] O\_CREAT and O\_EXCL are set and the named shared memory object already  
 41304 exists.

41305 [EINTR] The *shm\_open()* operation was interrupted by a signal.

41306 [EINVAL] The *shm\_open()* operation is not supported for the given name.

41307 [EMFILE] Too many file descriptors are currently in use by this process.

41308 [ENAMETOOLONG]

41309 The length of the *name* argument exceeds {PATH\_MAX} or a pathname  
 41310 component is longer than {NAME\_MAX}. |

41311 [ENFILE] Too many shared memory objects are currently open in the system.

41312 [ENOENT] O\_CREAT is not set and the named shared memory object does not exist.

41313 [ENOSPC] There is insufficient space for the creation of the new shared memory object.

#### 41314 EXAMPLES

41315 None.

#### 41316 APPLICATION USAGE

41317 None.

#### 41318 RATIONALE

41319 When the Memory Mapped Files option is supported, the normal *open()* call is used to obtain a  
 41320 descriptor to a file to be mapped according to existing practice with *mmap()*. When the Shared  
 41321 Memory Objects option is supported, the *shm\_open()* function shall obtain a descriptor to the  
 41322 shared memory object to be mapped. |

41323 There is ample precedent for having a file descriptor represent several types of objects. In the  
 41324 POSIX.1-1990 standard, a file descriptor can represent a file, a pipe, a FIFO, a tty, or a directory.  
 41325 Many implementations simply have an operations vector, which is indexed by the file descriptor  
 41326 type and does very different operations. Note that in some cases the file descriptor passed to  
 41327 generic operations on file descriptors are returned by *open()* or *creat()* and in some cases  
 41328 returned by alternate functions, such as *pipe()*. The latter technique is used by *shm\_open()*.

41329 Note that such shared memory objects can actually be implemented as mapped files. In both  
41330 cases, the size can be set after the open using *ftruncate()*. The *shm\_open()* function itself does not  
41331 create a shared object of a specified size because this would duplicate an extant function that set  
41332 the size of an object referenced by a file descriptor.

41333 On implementations where memory objects are implemented using the existing file system, the  
41334 *shm\_open()* function may be implemented using a macro that invokes *open()*, and the  
41335 *shm\_unlink()* function may be implemented using a macro that invokes *unlink()*.

41336 For implementations without a permanent file system, the definition of the name of the memory  
41337 objects is allowed not to survive a system reboot. Note that this allows systems with a  
41338 permanent file system to implement memory objects as data structures internal to the  
41339 implementation as well.

41340 On implementations that choose to implement memory objects using memory directly, a  
41341 *shm\_open()* followed by a *ftruncate()* and *close()* can be used to preallocate a shared memory  
41342 area and to set the size of that preallocation. This may be necessary for systems without virtual  
41343 memory hardware support in order to ensure that the memory is contiguous.

41344 The set of valid open flags to *shm\_open()* was restricted to *O\_RDONLY*, *O\_RDWR*, *O\_CREAT*,  
41345 and *O\_TRUNC* because these could be easily implemented on most memory mapping systems.  
41346 This volume of IEEE Std 1003.1-200x is silent on the results if the implementation cannot supply  
41347 the requested file access because of implementation-defined reasons, including hardware ones.

41348 The error conditions [EACCES] and [ENOTSUP] are provided to inform the application that the  
41349 implementation cannot complete a request.

41350 [EACCES] indicates for implementation-defined reasons, probably hardware-related, that the  
41351 implementation cannot comply with a requested mode because it conflicts with another  
41352 requested mode. An example might be that an application desires to open a memory object two  
41353 times, mapping different areas with different access modes. If the implementation cannot map a  
41354 single area into a process space in two places, which would be required if different access modes  
41355 were required for the two areas, then the implementation may inform the application at the time  
41356 of the second open.

41357 [ENOTSUP] indicates for implementation-defined reasons, probably hardware-related, that the  
41358 implementation cannot comply with a requested mode at all. An example would be that the  
41359 hardware of the implementation cannot support write-only shared memory areas.

41360 On all implementations, it may be desirable to restrict the location of the memory objects to  
41361 specific file systems for performance (such as a RAM disk) or implementation-defined reasons  
41362 (shared memory supported directly only on certain file systems). The *shm\_open()* function may  
41363 be used to enforce these restrictions. There are a number of methods available to the application  
41364 to determine an appropriate name of the file or the location of an appropriate directory. One  
41365 way is from the environment via *getenv()*. Another would be from a configuration file.

41366 This volume of IEEE Std 1003.1-200x specifies that memory objects have initial contents of zero  
41367 when created. This is consistent with current behavior for both files and newly allocated  
41368 memory. For those implementations that use physical memory, it would be possible that such  
41369 implementations could simply use available memory and give it to the process uninitialized.  
41370 This, however, is not consistent with standard behavior for the uninitialized data area, the stack,  
41371 and of course, files. Finally, it is highly desirable to set the allocated memory to zero for security  
41372 reasons. Thus, initializing memory objects to zero is required.

41373 **FUTURE DIRECTIONS**

41374 None.

41375 **SEE ALSO**

41376 *close()*, *dup()*, *exec*, *fcntl()*, *mmap()*, *shmat()*, *shmctl()*, *shmdt()*, *shm\_unlink()*, *umask()*, the Base  
41377 Definitions volume of IEEE Std 1003.1-200x, <fcntl.h>, <sys/mman.h>

41378 **CHANGE HISTORY**

41379 First released in Issue 5. Included for alignment with the POSIX Realtime Extension.

41380 **Issue 6**

41381 The *shm\_open()* function is marked as part of the Shared Memory Objects option.

41382 The [ENOSYS] error condition has been removed as stubs need not be provided if an  
41383 implementation does not support the Shared Memory Objects option.



41384 **NAME**41385 shm\_unlink — remove a shared memory object (**REALTIME**)41386 **SYNOPSIS**

41387 SHM #include &lt;sys/mman.h&gt;

41388 int shm\_unlink(const char \*name);

41389

41390 **DESCRIPTION**41391 The *shm\_unlink()* function shall remove the name of the shared memory object named by the  
41392 string pointed to by *name*.41393 If one or more references to the shared memory object exist when the object is unlinked, the  
41394 name shall be removed before *shm\_unlink()* returns, but the removal of the memory object  
41395 contents shall be postponed until all open and map references to the shared memory object have  
41396 been removed.41397 Even if the object continues to exist after the last *shm\_unlink()*, reuse of the name shall  
41398 subsequently cause *shm\_open()* to behave as if no shared memory object of this name exists (that  
41399 is, *shm\_open()* will fail if O\_CREAT is not set, or will create a new shared memory object if  
41400 O\_CREAT is set).41401 **RETURN VALUE**41402 Upon successful completion, a value of zero shall be returned. Otherwise, a value of -1 shall be  
41403 returned and *errno* set to indicate the error. If -1 is returned, the named shared memory object  
41404 shall not be changed by this function call.41405 **ERRORS**41406 The *shm\_unlink()* function shall fail if:

41407 [EACCES] Permission is denied to unlink the named shared memory object.

41408 [ENAMETOOLONG]

41409 The length of the *name* argument exceeds {PATH\_MAX} or a pathname  
41410 component is longer than {NAME\_MAX}.

41411 [ENOENT] The named shared memory object does not exist.

41412 **EXAMPLES**

41413 None.

41414 **APPLICATION USAGE**41415 Names of memory objects that were allocated with *open()* are deleted with *unlink()* in the usual  
41416 fashion. Names of memory objects that were allocated with *shm\_open()* are deleted with  
41417 *shm\_unlink()*. Note that the actual memory object is not destroyed until the last close and  
41418 unmap on it have occurred if it was already in use.41419 **RATIONALE**

41420 None.

41421 **FUTURE DIRECTIONS**

41422 None.

41423 **SEE ALSO**41424 *close()*, *mmap()*, *munmap()*, *shmat()*, *shmctl()*, *shmdt()*, *shm\_open()*, the Base Definitions volume  
41425 of IEEE Std 1003.1-200x, <sys/mman.h>

## 41426 CHANGE HISTORY

41427 First released in Issue 5. Included for alignment with the POSIX Realtime Extension.

## 41428 Issue 6

41429 The *shm\_unlink()* function is marked as part of the Shared Memory Objects option.

41430 In the DESCRIPTION, text is added to clarify that reusing the same name after a *shm\_unlink()*  
41431 will not attach to the old shared memory object.

41432 The [ENOSYS] error condition has been removed as stubs need not be provided if an  
41433 implementation does not support the Shared Memory Objects option.

41434 **NAME**41435 `shmat` — XSI shared memory attach operation41436 **SYNOPSIS**41437 XSI 

```
#include <sys/shm.h>
```

41438 

```
void *shmat(int shmid, const void *shmaddr, int shmflg);
```

41439

41440 **DESCRIPTION**

41441 The `shmat()` function operates on XSI shared memory (see the Base Definitions volume of  
 41442 IEEE Std 1003.1-200x, Section 3.340, Shared Memory Object). It is unspecified whether this  
 41443 function interoperates with the realtime interprocess communication facilities defined in Section  
 41444 2.8 (on page 491).

41445 The `shmat()` function attaches the shared memory segment associated with the shared memory  
 41446 identifier specified by `shmid` to the address space of the calling process. The segment is attached  
 41447 at the address specified by one of the following criteria:

- 41448 • If `shmaddr` is a null pointer, the segment is attached at the first available address as selected  
 41449 by the system.
- 41450 • If `shmaddr` is not a null pointer and `(shmflg & SHM_RND)` is non-zero, the segment is attached  
 41451 at the address given by `(shmaddr - ((uintptr_t)shmaddr % SHMLBA))`. The character `'%'` is the  
 41452 C-language remainder operator.
- 41453 • If `shmaddr` is not a null pointer and `(shmflg & SHM_RND)` is 0, the segment is attached at the  
 41454 address given by `shmaddr`.
- 41455 • The segment is attached for reading if `(shmflg & SHM_RDONLY)` is non-zero and the calling  
 41456 process has read permission; otherwise, if it is 0 and the calling process has read and write  
 41457 permission, the segment is attached for reading and writing.

41458 **RETURN VALUE**

41459 Upon successful completion, `shmat()` shall increment the value of `shm_nattch` in the data  
 41460 structure associated with the shared memory ID of the attached shared memory segment and  
 41461 return the segment's start address.

41462 Otherwise, the shared memory segment shall not be attached, `shmat()` shall return `-1`, and `errno`  
 41463 shall be set to indicate the error.

41464 **ERRORS**41465 The `shmat()` function shall fail if:

- |   |          |   |
|---|----------|---|
| 41466<br>41467                            | [EACCES] | Operation permission is denied to the calling process; see Section 2.7 (on page 489).   |
| 41468<br>41469<br>41470<br>41471<br>41472 | [EINVAL] | The value of <code>shmid</code> is not a valid shared memory identifier, the <code>shmaddr</code> is not a null pointer, and the value of <code>(shmaddr - ((uintptr_t)shmaddr % SHMLBA))</code> is an illegal address for attaching shared memory; or the <code>shmaddr</code> is not a null pointer, <code>(shmflg &amp; SHM_RND)</code> is 0, and the value of <code>shmaddr</code> is an illegal address for attaching shared memory. |
| 41473<br>41474                            | [EMFILE] | The number of shared memory segments attached to the calling process would exceed the system-imposed limit.   |
| 41475<br>41476                            | [ENOMEM] | The available data space is not large enough to accommodate the shared memory segment.  |

41477 **EXAMPLES**

41478 None.

41479 **APPLICATION USAGE**

41480 The POSIX Realtime Extension defines alternative interfaces for interprocess communication.  
41481 Application developers who need to use IPC should design their applications so that modules  
41482 using the IPC routines described in Section 2.7 (on page 489) can be easily modified to use the  
41483 alternative interfaces.

41484 **RATIONALE**

41485 None.

41486 **FUTURE DIRECTIONS**

41487 None.

41488 **SEE ALSO**

41489 *exec*, *exit()*, *fork()*, *shmctl()*, *shmdt()*, *shmget()*, *shm\_open()*, *shm\_unlink()*, the Base Definitions  
41490 volume of IEEE Std 1003.1-200x, <sys/shm.h>, Section 2.7 (on page 489)

41491 **CHANGE HISTORY**

41492 First released in Issue 2. Derived from Issue 2 of the SVID.

41493 **Issue 5**

41494 Moved from SHARED MEMORY to BASE.

41495 The note about use of POSIX Realtime Extension IPC routines has been moved from FUTURE  
41496 DIRECTIONS to a new APPLICATION USAGE section.

41497 **Issue 6**

41498 The Open Group Corrigendum U021/13 is applied.

41499 **NAME**

41500 shmctl — XSI shared memory control operations

41501 **SYNOPSIS**41502 XSI 

```
#include <sys/shm.h>
```

41503 

```
int shmctl(int shmid, int cmd, struct shmids *buf);
```

41504

41505 **DESCRIPTION**

41506 The *shmctl()* function operates on XSI shared memory (see the Base Definitions volume of  
 41507 IEEE Std 1003.1-200x, Section 3.340, Shared Memory Object). It is unspecified whether this  
 41508 function interoperates with the realtime interprocess communication facilities defined in Section  
 41509 2.8 (on page 491).

41510 The *shmctl()* function provides a variety of shared memory control operations as specified by  
 41511 *cmd*. The following values for *cmd* are available:

41512 **IPC\_STAT** Place the current value of each member of the **shmids** data structure  
 41513 associated with *shmid* into the structure pointed to by *buf*. The contents of the  
 41514 structure are defined in **<sys/shm.h>**.

41515 **IPC\_SET** Set the value of the following members of the **shmids** data structure  
 41516 associated with *shmid* to the corresponding value found in the structure  
 41517 pointed to by *buf*:

41518 shm\_perm.uid  
 41519 shm\_perm.gid  
 41520 shm\_perm.mode Low-order nine bits.

41521 **IPC\_SET** can only be executed by a process that has an effective user ID equal  
 41522 to either that of a process with appropriate privileges or to the value of  
 41523 *shm\_perm.cuid* or *shm\_perm.uid* in the **shmids** data structure associated with  
 41524 *shmid*.

41525 **IPC\_RMID** Remove the shared memory identifier specified by *shmid* from the system and  
 41526 destroy the shared memory segment and **shmids** data structure associated  
 41527 with it. **IPC\_RMID** can only be executed by a process that has an effective user  
 41528 ID equal to either that of a process with appropriate privileges or to the value  
 41529 of *shm\_perm.cuid* or *shm\_perm.uid* in the **shmids** data structure associated  
 41530 with *shmid*.

41531 **RETURN VALUE**

41532 Upon successful completion, *shmctl()* shall return 0; otherwise, it shall return -1 and set *errno* to  
 41533 indicate the error.

41534 **ERRORS**41535 The *shmctl()* function shall fail if:

41536 **[EACCES]** The argument *cmd* is equal to **IPC\_STAT** and the calling process does not have  
 41537 read permission; see Section 2.7 (on page 489).

41538 **[EINVAL]** The value of *shmid* is not a valid shared memory identifier, or the value of *cmd*  
 41539 is not a valid command.

41540 **[EPERM]** The argument *cmd* is equal to **IPC\_RMID** or **IPC\_SET** and the effective user ID  
 41541 of the calling process is not equal to that of a process with appropriate  
 41542 privileges and it is not equal to the value of *shm\_perm.cuid* or *shm\_perm.uid* in  
 41543 the data structure associated with *shmid*.

41544 The *shmctl()* function may fail if:

41545 [EOVERFLOW] The *cmd* argument is `IPC_STAT` and the *gid* or *uid* value is too large to be  
41546 stored in the structure pointed to by the *buf* argument.

41547 **EXAMPLES**

41548 None.

41549 **APPLICATION USAGE**

41550 The POSIX Realtime Extension defines alternative interfaces for interprocess communication.  
41551 Application developers who need to use IPC should design their applications so that modules  
41552 using the IPC routines described in Section 2.7 (on page 489) can be easily modified to use the  
41553 alternative interfaces.

41554 **RATIONALE**

41555 None.

41556 **FUTURE DIRECTIONS**

41557 None.

41558 **SEE ALSO**

41559 *shmat()*, *shmdt()*, *shmget()*, *shm\_open()*, *shm\_unlink()*, the Base Definitions volume of  
41560 IEEE Std 1003.1-200x, <sys/shm.h>, Section 2.7 (on page 489)

41561 **CHANGE HISTORY**

41562 First released in Issue 2. Derived from Issue 2 of the SVID.

41563 **Issue 5**

41564 Moved from SHARED MEMORY to BASE.

41565 The note about use of POSIX Realtime Extension IPC routines has been moved from FUTURE  
41566 DIRECTIONS to a new APPLICATION USAGE section.

41567 **NAME**

41568       shmdt — XSI shared memory detach operation

41569 **SYNOPSIS**

41570 XSI       #include &lt;sys/shm.h&gt;

41571       int shmdt(const void \*shmaddr);

41572

41573 **DESCRIPTION**

41574       The *shmdt()* function operates on XSI shared memory (see the Base Definitions volume of  
 41575       IEEE Std 1003.1-200x, Section 3.340, Shared Memory Object). It is unspecified whether this  
 41576       function interoperates with the realtime interprocess communication facilities defined in Section  
 41577       2.8 (on page 491).

41578       The *shmdt()* function detaches the shared memory segment located at the address specified by  
 41579       *shmaddr* from the address space of the calling process.

41580 **RETURN VALUE**

41581       Upon successful completion, *shmdt()* shall decrement the value of *shm\_nattch* in the data  
 41582       structure associated with the shared memory ID of the attached shared memory segment and  
 41583       return 0.

41584       Otherwise, the shared memory segment shall not be detached, *shmdt()* shall return  $-1$ , and *errno*  
 41585       shall be set to indicate the error.

41586 **ERRORS**41587       The *shmdt()* function shall fail if:

41588       [EINVAL]       The value of *shmaddr* is not the data segment start address of a shared  
 41589       memory segment.

41590 **EXAMPLES**

41591       None.

41592 **APPLICATION USAGE**

41593       The POSIX Realtime Extension defines alternative interfaces for interprocess communication.  
 41594       Application developers who need to use IPC should design their applications so that modules  
 41595       using the IPC routines described in Section 2.7 (on page 489) can be easily modified to use the  
 41596       alternative interfaces.

41597 **RATIONALE**

41598       None.

41599 **FUTURE DIRECTIONS**

41600       None.

41601 **SEE ALSO**

41602       *exec*, *exit()*, *fork()*, *shmat()*, *shmctl()*, *shmget()*, *shm\_open()*, *shm\_unlink()*, the Base Definitions  
 41603       volume of IEEE Std 1003.1-200x, <sys/shm.h>, Section 2.7 (on page 489)

41604 **CHANGE HISTORY**

41605       First released in Issue 2. Derived from Issue 2 of the SVID.

41606 **Issue 5**

41607       Moved from SHARED MEMORY to BASE.

41608       The note about use of POSIX Realtime Extension IPC routines has been moved from FUTURE  
 41609       DIRECTIONS to a new APPLICATION USAGE section.

## 41610 NAME

41611 shmget — get XSI shared memory segment

## 41612 SYNOPSIS

41613 XSI 

```
#include <sys/shm.h>
```

41614 

```
int shmget(key_t key, size_t size, int shmflg);
```

41615

## 41616 DESCRIPTION

41617 The *shmget()* function operates on XSI shared memory (see the Base Definitions volume of  
 41618 IEEE Std 1003.1-200x, Section 3.340, Shared Memory Object). It is unspecified whether this  
 41619 function interoperates with the realtime interprocess communication facilities defined in Section  
 41620 2.8 (on page 491).

41621 The *shmget()* function shall return the shared memory identifier associated with *key*.

41622 A shared memory identifier, associated data structure, and shared memory segment of at least  
 41623 *size* bytes (see <sys/shm.h>) are created for *key* if one of the following is true:

- 41624 • The argument *key* is equal to `IPC_PRIVATE`.
- 41625 • The argument *key* does not already have a shared memory identifier associated with it and  
 41626 (*shmflg* & `IPC_CREAT`) is non-zero.

41627 Upon creation, the data structure associated with the new shared memory identifier shall be  
 41628 initialized as follows:

- 41629 • The values of *shm\_perm.cuid*, *shm\_perm.uid*, *shm\_perm.cgid*, and *shm\_perm.gid* are set equal to  
 41630 the effective user ID and effective group ID, respectively, of the calling process.
- 41631 • The low-order nine bits of *shm\_perm.mode* are set equal to the low-order nine bits of *shmflg*.  
 41632 The value of *shm\_segsz* is set equal to the value of *size*.
- 41633 • The values of *shm\_lpid*, *shm\_nattch*, *shm\_atime*, and *shm\_dtime* are set equal to 0.
- 41634 • The value of *shm\_ctime* is set equal to the current time.

41635 When the shared memory segment is created, it shall be initialized with all zero values.

## 41636 RETURN VALUE

41637 Upon successful completion, *shmget()* shall return a non-negative integer, namely a shared  
 41638 memory identifier; otherwise, it shall return `-1` and set *errno* to indicate the error.

## 41639 ERRORS

41640 The *shmget()* function shall fail if:

- |       |          |   |
|-------|----------|---|
| 41641 | [EACCES] | A shared memory identifier exists for <i>key</i> but operation permission as<br>41642 specified by the low-order nine bits of <i>shmflg</i> would not be granted; see<br>41643 Section 2.7 (on page 489).       |
| 41644 | [EEXIST] | A shared memory identifier exists for the argument <i>key</i> but ( <i>shmflg</i><br>41645 & <code>IPC_CREAT</code> ) && ( <i>shmflg</i> & <code>IPC_EXCL</code> ) is non-zero.                                 |
| 41646 | [EINVAL] | A shared memory segment is to be created and the value of <i>size</i> is less than<br>41647 the system-imposed minimum or greater than the system-imposed maximum.  |
| 41648 | [EINVAL] | No shared memory segment is to be created and a shared memory segment<br>41649 exists for <i>key</i> but the size of the segment associated with it is less than <i>size</i> and<br>41650 <i>size</i> is not 0. |



- 41651 [ENOENT] A shared memory identifier does not exist for the argument *key* and (*shmflg*  
41652 &IPC\_CREAT) is 0.
- 41653 [ENOMEM] A shared memory identifier and associated shared memory segment shall be  
41654 created, but the amount of available physical memory is not sufficient to fill  
41655 the request.
- 41656 [ENOSPC] A shared memory identifier is to be created, but the system-imposed limit on  
41657 the maximum number of allowed shared memory identifiers system-wide  
41658 would be exceeded.

**41659 EXAMPLES**

41660 None.

**41661 APPLICATION USAGE**

41662 The POSIX Realtime Extension defines alternative interfaces for interprocess communication.  
41663 Application developers who need to use IPC should design their applications so that modules  
41664 using the IPC routines described in Section 2.7 (on page 489) can be easily modified to use the  
41665 alternative interfaces.

**41666 RATIONALE**

41667 None.

**41668 FUTURE DIRECTIONS**

41669 None.

**41670 SEE ALSO**

41671 *shmat()*, *shmctl()*, *shmdt()*, *shm\_open()*, *shm\_unlink()*, the Base Definitions volume of  
41672 IEEE Std 1003.1-200x, <sys/shm.h>, Section 2.7 (on page 489)

**41673 CHANGE HISTORY**

41674 First released in Issue 2. Derived from Issue 2 of the SVID.

**41675 Issue 5**

41676 Moved from SHARED MEMORY to BASE.

41677 The note about use of POSIX Realtime Extension IPC routines has been moved from FUTURE  
41678 DIRECTIONS to a new APPLICATION USAGE section.

41679 **NAME**

41680 shutdown — shut down socket send and receive operations

41681 **SYNOPSIS**

41682 #include <sys/socket.h>

41683 int shutdown(int *socket*, int *how*);

41684 **DESCRIPTION**

41685 The *shutdown()* function shall cause all or part of a full-duplex connection on the socket  
41686 associated with the file descriptor *socket* to be shut down.

41687 The *shutdown()* function takes the following arguments:

- |       |               |  |
|-------|---------------|--|
| 41688 | <i>socket</i> | Specifies the file descriptor of the socket.               |
| 41689 | <i>how</i>    | Specifies the type of shutdown. The values are as follows: |
| 41690 | SHUT_RD       | Disables further receive operations.                       |
| 41691 | SHUT_WR       | Disables further send operations.                          |
| 41692 | SHUT_RDWR     | Disables further send and receive operations.              |

41693 The *shutdown()* function disables subsequent send and/or receive operations on a socket,  
41694 depending on the value of the *how* argument.

41695 **RETURN VALUE**

41696 Upon successful completion, *shutdown()* shall return 0; otherwise, -1 shall be returned and *errno*  
41697 set to indicate the error.

41698 **ERRORS**

41699 The *shutdown()* function shall fail if:

- |       |            |  |
|-------|------------|--|
| 41700 | [EBADF]    | The <i>socket</i> argument is not a valid file descriptor. |
| 41701 | [EINVAL]   | The <i>how</i> argument is invalid.                        |
| 41702 | [ENOTCONN] | The socket is not connected.                               |
| 41703 | [ENOTSOCK] | The <i>socket</i> argument does not refer to a socket.     |
- 41704 The *shutdown()* function may fail if:
- |       |           |   |
|-------|-----------|---|
| 41705 | [ENOBUFS] | Insufficient resources were available in the system to perform the operation. |
|-------|-----------|---|

41706 **EXAMPLES**

41707 None.

41708 **APPLICATION USAGE**

41709 None.

41710 **RATIONALE**

41711 None.

41712 **FUTURE DIRECTIONS**

41713 None.

41714 **SEE ALSO**

41715 *getsockopt()*, *read()*, *recv()*, *recvfrom()*, *recvmsg()*, *select()*, *send()*, *sendto()*, *setsockopt()*, *socket()*,  
41716 *write()*, the Base Definitions volume of IEEE Std 1003.1-200x, <sys/socket.h>

41717 **CHANGE HISTORY**

41718 First released in Issue 6. Derived from the XNS, Issue 5.2 specification.

## 41719 NAME

41720 sigaction — examine and change signal action

## 41721 SYNOPSIS

41722 cx #include &lt;signal.h&gt;

41723 int sigaction(int sig, const struct sigaction \*restrict act,  
41724 struct sigaction \*restrict oact);

41725

## 41726 DESCRIPTION

41727 The *sigaction()* function allows the calling process to examine and/or specify the action to be  
41728 associated with a specific signal. The argument *sig* specifies the signal; acceptable values are  
41729 defined in <signal.h>.41730 The structure **sigaction**, used to describe an action to be taken, is defined in the <signal.h>  
41731 header to include at least the following members:

41732

41733

41734

41735

41736

41737

41738

41739

41740

Member Type	Member Name	Description
void(*) (int) sigset_t	sa_handler sa_mask	SIG_DFL, SIG_IGN, or pointer to a function. Additional set of signals to be blocked during execution of signal-catching function.
int	sa_flags	Special flags to affect behavior of signal.
void(*) (int, siginfo_t *, void *)	sa_sigaction	Signal-catching function.

41741 The storage occupied by *sa\_handler* and *sa\_sigaction* may overlap, and a conforming application  
41742 shall not use both simultaneously.41743 If the argument *act* is not a null pointer, it points to a structure specifying the action to be  
41744 associated with the specified signal. If the argument *oact* is not a null pointer, the action  
41745 previously associated with the signal is stored in the location pointed to by the argument *oact*. If  
41746 the argument *act* is a null pointer, signal handling is unchanged; thus, the call can be used to  
41747 enquire about the current handling of a given signal. The SIGKILL and SIGSTOP signals shall  
41748 not be added to the signal mask using this mechanism; this restriction shall be enforced by the  
41749 system without causing an error to be indicated.41750 If the SA\_SIGINFO flag (see below) is cleared in the *sa\_flags* field of the **sigaction** structure, the  
41751 XSI|RTS *sa\_handler* field identifies the action to be associated with the specified signal. If the  
41752 SA\_SIGINFO flag is set in the *sa\_flags* field, and the implementation supports the Realtime  
41753 Signals Extension option or the X/Open System Interfaces Extension option, the *sa\_sigaction*  
41754 field specifies a signal-catching function. If the SA\_SIGINFO bit is cleared and the *sa\_handler*  
41755 field specifies a signal-catching function, or if the SA\_SIGINFO bit is set, the *sa\_mask* field  
41756 identifies a set of signals that shall be added to the signal mask of the thread before the signal-  
41757 catching function is invoked. If the *sa\_handler* field specifies a signal-catching function, the  
41758 *sa\_mask* field identifies a set of signals that shall be added to the process' signal mask before the  
41759 signal-catching function is invoked.41760 The *sa\_flags* field can be used to modify the behavior of the specified signal.41761 The following flags, defined in the <signal.h> header, can be set in *sa\_flags*:41762 XSI SA\_NOCLDSTOP Do not generate SIGCHLD when children stop or stopped children  
41763 continue.

41764		If <i>sig</i> is SIGCHLD and the SA_NOCLDSTOP flag is not set in <i>sa_flags</i> , and the implementation supports the SIGCHLD signal, then a SIGCHLD signal shall be generated for the calling process whenever any of its child processes stop and a SIGCHLD signal may be generated for the calling process whenever any of its stopped child processes are continued. If <i>sig</i> is SIGCHLD and the SA_NOCLDSTOP flag is set in <i>sa_flags</i> , then the implementation shall not generate a SIGCHLD signal in this way.
41765		
41766		
41767 XSI		
41768		
41769		
41770		
41771 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.
41772		
41773		
41774 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.
41775		
41776		<b>Note:</b> SIGILL and SIGTRAP cannot be automatically reset when delivered; the system silently enforces this restriction.
41777		
41778		Otherwise, the disposition of the signal shall not be modified on entry to the signal handler.
41779		
41780		In addition, if this flag is set, <i>sigaction()</i> behaves as if the SA_NODEFER flag were also set.
41781		
41782 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].
41783		
41784		
41785		
41786		
41787		
41788	SA_SIGINFO	If cleared and the signal is caught, the signal-catching function shall be entered as:
41789		
41790		<pre>void func(int signo);</pre>
41791		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.
41792		
41793		
41794		
41795 XSI RTS		If SA_SIGINFO is set and the signal is caught, the signal-catching function shall be entered as:
41796		
41797		<pre>void func(int signo, siginfo_t *info, void *context);</pre>
41798		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.
41799		
41800		
41801		
41802		
41803		
41804		
41805		
41806		The <i>si_signo</i> member contains the system-generated signal number.
41807 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.
41808		
41809		

41810 XSI|RTS           The *si\_code* member contains a code identifying the cause of the signal.

41811 XSI               If the value of *si\_code* is less than or equal to 0, then the signal was  
41812                       generated by a process and *si\_pid* and *si\_uid*, respectively, indicate the  
41813                       process ID and the real user ID of the sender. The `<signal.h>` header  
41814                       description contains information about the signal specific contents of the  
41815                       elements of the **siginfo\_t** type.

41816 XSI           SA\_NOCLDWAIT   If set, and *sig* equals SIGCHLD, child processes of the calling processes  
41817                       shall not be transformed into zombie processes when they terminate. If  
41818                       the calling process subsequently waits for its children, and the process  
41819                       has no unwaited-for children that were transformed into zombie  
41820                       processes, it shall block until all of its children terminate, and *wait()*,  
41821                       *waitid()*, and *waitpid()* shall fail and set *errno* to [ECHILD]. Otherwise,  
41822                       terminating child processes shall be transformed into zombie processes,  
41823                       unless SIGCHLD is set to SIG\_IGN.

41824 XSI           SA\_NODEFER       If set and *sig* is caught, *sig* shall not be added to the process' signal mask  
41825                       on entry to the signal handler unless it is included in *sa\_mask*. Otherwise,  
41826                       *sig* shall always be added to the process' signal mask on entry to the  
41827                       signal handler.

41828                       When a signal is caught by a signal-catching function installed by *sigaction()*, a new signal mask  
41829                       is calculated and installed for the duration of the signal-catching function (or until a call to either  
41830                       *sigprocmask()* or *sigsuspend()* is made). This mask is formed by taking the union of the current  
41831 XSI               signal mask and the value of the *sa\_mask* for the signal being delivered unless SA\_NODEFER or  
41832                       SA\_RESETHAND is set, and then including the signal being delivered. If and when the user's  
41833                       signal handler returns normally, the original signal mask is restored.

41834                       Once an action is installed for a specific signal, it shall remain installed until another action is  
41835 XSI               explicitly requested (by another call to *sigaction()*), until the SA\_RESETHAND flag causes  
41836                       resetting of the handler, or until one of the *exec* functions is called.

41837                       If the previous action for *sig* had been established by *signal()*, the values of the fields returned in  
41838                       the structure pointed to by *oact* are unspecified, and in particular *oact->sa\_handler* is not  
41839                       necessarily the same value passed to *signal()*. However, if a pointer to the same structure or a  
41840                       copy thereof is passed to a subsequent call to *sigaction()* via the *act* argument, handling of the  
41841                       signal shall be as if the original call to *signal()* were repeated.

41842                       If *sigaction()* fails, no new signal handler is installed.

41843                       It is unspecified whether an attempt to set the action for a signal that cannot be caught or  
41844                       ignored to SIG\_DFL is ignored or causes an error to be returned with *errno* set to [EINVAL].

41845                       If SA\_SIGINFO is not set in *sa\_flags*, then the disposition of subsequent occurrences of *sig* when  
41846                       it is already pending is implementation-defined; the signal-catching function shall be invoked  
41847 RTS               with a single argument. If the implementation supports the Realtime Signals Extension option,  
41848                       and if SA\_SIGINFO is set in *sa\_flags*, then subsequent occurrences of *sig* generated by *sigqueue()*  
41849                       or as a result of any signal-generating function that supports the specification of an application-  
41850                       defined value (when *sig* is already pending) shall be queued in FIFO order until delivered or  
41851                       accepted; the signal-catching function shall be invoked with three arguments. The application  
41852                       specified value is passed to the signal-catching function as the *si\_value* member of the **siginfo\_t**  
41853                       structure.

41854                       The result of the use of *sigaction()* and a *sigwait()* function concurrently within a process on the  
41855                       same signal is unspecified.

## 41856 RETURN VALUE

41857 Upon successful completion, *sigaction()* shall return 0; otherwise, -1 shall be returned, *errno* shall  
 41858 be set to indicate the error, and no new signal-catching function shall be installed.

## 41859 ERRORS

41860 The *sigaction()* function shall fail if:

41861 [EINVAL] The *sig* argument is not a valid signal number or an attempt is made to catch a  
 41862 signal that cannot be caught or ignore a signal that cannot be ignored.

41863 [ENOTSUP] The SA\_SIGINFO bit flag is set in the *sa\_flags* field of the **sigaction** structure,  
 41864 and the implementation does not support either the Realtime Signals  
 41865 Extension option, or the X/Open System Interfaces Extension option.

41866 The *sigaction()* function may fail if:

41867 [EINVAL] An attempt was made to set the action to SIG\_DFL for a signal that cannot be  
 41868 caught or ignored (or both).

## 41869 EXAMPLES

41870 None.

## 41871 APPLICATION USAGE

41872 The *sigaction()* function supersedes the *signal()* function, and should be used in preference. In  
 41873 particular, *sigaction()* and *signal()* should not be used in the same process to control the same  
 41874 signal. The behavior of reentrant functions, as defined in the DESCRIPTION, is as specified by  
 41875 this volume of IEEE Std 1003.1-200x, regardless of invocation from a signal-catching function.  
 41876 This is the only intended meaning of the statement that reentrant functions may be used in  
 41877 signal-catching functions without restrictions. Applications must still consider all effects of such  
 41878 functions on such things as data structures, files, and process state. In particular, application  
 41879 writers need to consider the restrictions on interactions when interrupting *sleep()* and  
 41880 interactions among multiple handles for a file description. The fact that any specific function is  
 41881 listed as reentrant does not necessarily mean that invocation of that function from a signal-  
 41882 catching function is recommended.

41883 In order to prevent errors arising from interrupting non-reentrant function calls, applications  
 41884 should protect calls to these functions either by blocking the appropriate signals or through the  
 41885 use of some programmatic semaphore (see *semget()*, *sem\_init()*, *sem\_open()*, and so on). Note in  
 41886 particular that even the “safe” functions may modify *errno*; the signal-catching function, if not  
 41887 executing as an independent thread, may want to save and restore its value. Naturally, the same  
 41888 principles apply to the reentrancy of application routines and asynchronous data access. Note  
 41889 that *longjmp()* and *siglongjmp()* are not in the list of reentrant functions. This is because the code  
 41890 executing after *longjmp()* and *siglongjmp()* can call any unsafe functions with the same danger as  
 41891 calling those unsafe functions directly from the signal handler. Applications that use *longjmp()*  
 41892 and *siglongjmp()* from within signal handlers require rigorous protection in order to be portable.  
 41893 Many of the other functions that are excluded from the list are traditionally implemented using  
 41894 either *malloc()* or *free()* functions or the standard I/O library, both of which traditionally use  
 41895 data structures in a non-reentrant manner. Since any combination of different functions using a  
 41896 common data structure can cause reentrancy problems, this volume of IEEE Std 1003.1-200x  
 41897 does not define the behavior when any unsafe function is called in a signal handler that  
 41898 interrupts an unsafe function.

41899 If the signal occurs other than as the result of calling *abort()*, *kill()*, or *raise()*, the behavior is  
 41900 undefined if the signal handler calls any function in the standard library other than one of the  
 41901 functions listed in the table above or refers to any object with static storage duration other than  
 41902 by assigning a value to a static storage duration variable of type **volatile sig\_atomic\_t**.  
 41903 Furthermore, if such a call fails, the value of *errno* is unspecified.

41904 Usually, the signal is executed on the stack that was in effect before the signal was delivered. An  
41905 alternate stack may be specified to receive a subset of the signals being caught.

41906 When the signal handler returns, the receiving process resumes execution at the point it was  
41907 interrupted unless the signal handler makes other arrangements. If *longjmp()* or *\_longjmp()*  
41908 is used to leave the signal handler, then the signal mask must be explicitly restored by the process.

41909 This volume of IEEE Std 1003.1-200x defines the third argument of a signal handling function  
41910 when SA\_SIGINFO is set as a **void** \* instead of a **ucontext\_t** \*, but without requiring type  
41911 checking. New applications should explicitly cast the third argument of the signal handling  
41912 function to **ucontext\_t** \*.

41913 The BSD optional four argument signal handling function is not supported by this volume of  
41914 IEEE Std 1003.1-200x. The BSD declaration would be:

```
41915 void handler(int sig, int code, struct sigcontext *scp,  
41916             char *addr);
```

41917 where *sig* is the signal number, *code* is additional information on certain signals, *scp* is a pointer  
41918 to the sigcontext structure, and *addr* is additional address information. Much the same  
41919 information is available in the objects pointed to by the second argument of the signal handler  
41920 specified when SA\_SIGINFO is set.

#### 41921 RATIONALE

41922 Although this volume of IEEE Std 1003.1-200x requires that signals that cannot be ignored shall  
41923 not be added to the signal mask when a signal-catching function is entered, there is no explicit  
41924 requirement that subsequent calls to *sigaction()* reflect this in the information returned in the *oact*  
41925 argument. In other words, if SIGKILL is included in the *sa\_mask* field of *act*, it is unspecified  
41926 whether or not a subsequent call to *sigaction()* returns with SIGKILL included in the *sa\_mask*  
41927 field of *oact*.

41928 The SA\_NOCLDSTOP flag, when supplied in the *act->sa\_flags* parameter, allows overloading  
41929 SIGCHLD with the System V semantics that each SIGCLD signal indicates a single terminated |  
41930 child. Most conforming applications that catch SIGCHLD are expected to install signal-catching |  
41931 functions that repeatedly call the *waitpid()* function with the WNOHANG flag set, acting on  
41932 each child for which status is returned, until *waitpid()* returns zero. If stopped children are not of  
41933 interest, the use of the SA\_NOCLDSTOP flag can prevent the overhead from invoking the  
41934 signal-catching routine when they stop.

41935 Some historical implementations also define other mechanisms for stopping processes, such as  
41936 the *ptrace()* function. These implementations usually do not generate a SIGCHLD signal when  
41937 processes stop due to this mechanism; however, that is beyond the scope of this volume of  
41938 IEEE Std 1003.1-200x.

41939 This volume of IEEE Std 1003.1-200x requires that calls to *sigaction()* that supply a NULL *act*  
41940 argument succeed, even in the case of signals that cannot be caught or ignored (that is, SIGKILL  
41941 or SIGSTOP). The System V *signal()* and BSD *sigvec()* functions return [EINVAL] in these cases  
41942 and, in this respect, their behavior varies from *sigaction()*.

41943 This volume of IEEE Std 1003.1-200x requires that *sigaction()* properly save and restore a signal  
41944 action set up by the ISO C standard *signal()* function. However, there is no guarantee that the  
41945 reverse is true, nor could there be given the greater amount of information conveyed by the  
41946 **sigaction** structure. Because of this, applications should avoid using both functions for the same  
41947 signal in the same process. Since this cannot always be avoided in case of general-purpose  
41948 library routines, they should always be implemented with *sigaction()*.

41949 It was intended that the *signal()* function should be implementable as a library routine using  
41950 *sigaction()*.



41951 The POSIX Realtime Extension extends the *sigaction()* function as specified by the POSIX.1-1990  
 41952 standard to allow the application to request on a per-signal basis via an additional signal action  
 41953 flag that the extra parameters, including the application-defined signal value, if any, be passed  
 41954 to the signal-catching function.

#### 41955 FUTURE DIRECTIONS

41956 None.

#### 41957 SEE ALSO

41958 Section 2.4 (on page 478), *bsd\_signal()*, *kill()*, *\_longjmp()*, *longjmp()*, *raise()*, *semget()*, *sem\_init()*,  
 41959 *sem\_open()*, *sigaddset()*, *sigaltstack()*, *sigdelset()*, *sigemptyset()*, *sigfillset()*, *sigismember()*, *signal()*,  
 41960 *sigprocmask()*, *sigsuspend()*, *wait()*, *waitid()*, *waitpid()*, the Base Definitions volume of  
 41961 IEEE Std 1003.1-200x, <**signal.h**>, <**ucontext.h**>

#### 41962 CHANGE HISTORY

41963 First released in Issue 3.

41964 Entry included for alignment with the POSIX.1-1988 standard.

#### 41965 Issue 5

41966 The DESCRIPTION is updated for alignment with the POSIX Realtime Extension and POSIX  
 41967 Threads Extension.

41968 In the DESCRIPTION, the second argument to *func* when SA\_SIGINFO is set is no longer  
 41969 permitted to be NULL, and the description of permitted **siginfo\_t** contents is expanded by  
 41970 reference to <**signal.h**>.

41971 Since the X/OPEN UNIX Extension functionality is now folded into the BASE, the [ENOTSUP]  
 41972 error is deleted.

#### 41973 Issue 6

41974 The Open Group Corrigendum U028/7 is applied. In the paragraph entitled “Signal Effects on  
 41975 Other Functions”, a reference to *sigpending()* is added.

41976 In the DESCRIPTION, the text “Signal Generation and Delivery”, “Signal Actions”, and “Signal  
 41977 Effects on Other Functions” are moved to a separate section of this volume of  
 41978 IEEE Std 1003.1-200x.

41979 Text describing functionality from the Realtime Signals option is marked.

41980 The following changes are made for alignment with the ISO POSIX-1: 1996 standard:

- 41981 • The [ENOTSUP] error condition is added.

41982 The DESCRIPTION is updated to avoid use of the term “must” for application requirements.

41983 The **restrict** keyword is added to the *sigaction()* prototype for alignment with the  
 41984 ISO/IEC 9899: 1999 standard.

41985 References to the *wait3()* function are removed.

41986 The SYNOPSIS is marked CX since the presence of this function in the <**signal.h**> header is an  
 41987 extension over the ISO C standard.

41988 **NAME**

41989 sigaddset — add a signal to a signal set

41990 **SYNOPSIS**41991 **CX** #include <signal.h>

41992 int sigaddset(sigset\_t \*set, int signo);

41993

41994 **DESCRIPTION**41995 The *sigaddset()* function adds the individual signal specified by the *signo* to the signal set pointed  
41996 to by *set*.41997 Applications shall call either *sigemptyset()* or *sigfillset()* at least once for each object of type  
41998 **sigset\_t** prior to any other use of that object. If such an object is not initialized in this way, but is  
41999 nonetheless supplied as an argument to any of *pthread\_sigmask()*, *sigaction()*, *sigaddset()*,  
42000 *sigdelset()*, *sigismember()*, *sigpending()*, *sigprocmask()*, *sigsuspend()*, *sigtimedwait()*, *sigwait()*, or  
42001 *sigwaitinfo()*, the results are undefined.42002 **RETURN VALUE**42003 Upon successful completion, *sigaddset()* shall return 0; otherwise, it shall return -1 and set *errno*  
42004 to indicate the error.42005 **ERRORS**42006 The *sigaddset()* function may fail if:42007 [EINVAL] The value of the *signo* argument is an invalid or unsupported signal number.42008 **EXAMPLES**

42009 None.

42010 **APPLICATION USAGE**

42011 None.

42012 **RATIONALE**

42013 None.

42014 **FUTURE DIRECTIONS**

42015 None.

42016 **SEE ALSO**42017 Section 2.4 (on page 478), *sigaction()*, *sigdelset()*, *sigemptyset()*, *sigfillset()*, *sigismember()*,  
42018 *sigpending()*, *sigprocmask()*, *sigsuspend()*, the Base Definitions volume of IEEE Std 1003.1-200x,  
42019 <**signal.h**>42020 **CHANGE HISTORY**

42021 First released in Issue 3.

42022 Entry included for alignment with the POSIX.1-1988 standard.

42023 **Issue 5**42024 The last paragraph of the DESCRIPTION was included as an APPLICATION USAGE note in  
42025 previous issues.42026 **Issue 6**

42027 The DESCRIPTION is updated to avoid use of the term “must” for application requirements.

42028 The SYNOPSIS is marked CX since the presence of this function in the <**signal.h**> header is an  
42029 extension over the ISO C standard.

42030 **NAME**

42031 sigaltstack — set and get signal alternate stack context

42032 **SYNOPSIS**42033 XSI 

```
#include <signal.h>
```

42034 

```
int sigaltstack(const stack_t *restrict ss, stack_t *restrict oss);
```

42035

42036 **DESCRIPTION**

42037 The *sigaltstack()* function allows a process to define and examine the state of an alternate stack  
 42038 for signal handlers. Signals that have been explicitly declared to execute on the alternate stack  
 42039 shall be delivered on the alternate stack.

42040 If *ss* is not a null pointer, it points to a **stack\_t** structure that specifies the alternate signal stack  
 42041 that shall take effect upon return from *sigaltstack()*. The *ss\_flags* member specifies the new stack  
 42042 state. If it is set to `SS_DISABLE`, the stack is disabled and *ss\_sp* and *ss\_size* are ignored.  
 42043 Otherwise, the stack shall be enabled, and the *ss\_sp* and *ss\_size* members specify the new address  
 42044 and size of the stack.

42045 The range of addresses starting at *ss\_sp* up to but not including *ss\_sp+ss\_size*, is available to the  
 42046 implementation for use as the stack. This function makes no assumptions regarding which end  
 42047 is the stack base and in which direction the stack grows as items are pushed.

42048 If *oss* is not a null pointer, on successful completion it shall point to a **stack\_t** structure that  
 42049 specifies the alternate signal stack that was in effect prior to the call to *sigaltstack()*. The *ss\_sp*  
 42050 and *ss\_size* members specify the address and size of that stack. The *ss\_flags* member specifies the  
 42051 stack's state, and may contain one of the following values:

42052 **SS\_ONSTACK** The process is currently executing on the alternate signal stack. Attempts to  
 42053 modify the alternate signal stack while the process is executing on it fail. This  
 42054 flag shall not be modified by processes.

42055 **SS\_DISABLE** The alternate signal stack is currently disabled.

42056 The value `SIGSTKSZ` is a system default specifying the number of bytes that would be used to  
 42057 cover the usual case when manually allocating an alternate stack area. The value `MINSIGSTKSZ`  
 42058 is defined to be the minimum stack size for a signal handler. In computing an alternate stack  
 42059 size, a program should add that amount to its stack requirements to allow for the system  
 42060 implementation overhead. The constants `SS_ONSTACK`, `SS_DISABLE`, `SIGSTKSZ`, and  
 42061 `MINSIGSTKSZ` are defined in `<signal.h>`.

42062 After a successful call to one of the *exec* functions, there are no alternate signal stacks in the new  
 42063 process image.

42064 In some implementations, a signal (whether or not indicated to execute on the alternate stack)  
 42065 shall always execute on the alternate stack if it is delivered while another signal is being caught  
 42066 using the alternate stack.

42067 Use of this function by library threads that are not bound to kernel-scheduled entities results in  
 42068 undefined behavior.

42069 **RETURN VALUE**

42070 Upon successful completion, *sigaltstack()* shall return 0; otherwise, it shall return `-1` and set *errno*  
 42071 to indicate the error.

42072 **ERRORS**

42073 The *sigaltstack()* function shall fail if:

42074 [EINVAL] The *ss* argument is not a null pointer, and the *ss\_flags* member pointed to by *ss*  
42075 contains flags other than *SS\_DISABLE*.

42076 [ENOMEM] The size of the alternate stack area is less than *MINSIGSTKSZ*.

42077 [EPERM] An attempt was made to modify an active stack.

42078 **EXAMPLES**42079 **Allocating Memory for an Alternate Stack**

42080 The following example illustrates a method for allocating memory for an alternate stack.

```
42081 #include <signal.h>
42082 ...
42083 if ((sigstk.ss_sp = malloc(SIGSTKSZ)) == NULL)
42084     /* Error return. */
42085     sigstk.ss_size = SIGSTKSZ;
42086     sigstk.ss_flags = 0;
42087     if (sigaltstack(&sigstk, (stack_t *)0) < 0)
42088         perror("sigaltstack");
```

42089 **APPLICATION USAGE**

42090 On some implementations, stack space is automatically extended as needed. On those  
42091 implementations, automatic extension is typically not available for an alternate stack. If the stack  
42092 overflows, the behavior is undefined.

42093 **RATIONALE**

42094 None.

42095 **FUTURE DIRECTIONS**

42096 None.

42097 **SEE ALSO**

42098 Section 2.4 (on page 478), *sigaction()*, *sigsetjmp()*, the Base Definitions volume of  
42099 IEEE Std 1003.1-200x, <signal.h>

42100 **CHANGE HISTORY**

42101 First released in Issue 4, Version 2.

42102 **Issue 5**

42103 Moved from X/OPEN UNIX extension to BASE.

42104 The last sentence of the DESCRIPTION was included as an APPLICATION USAGE note in  
42105 previous issues.

42106 **Issue 6**

42107 The DESCRIPTION is updated to avoid use of the term “must” for application requirements.

42108 The **restrict** keyword is added to the *sigaltstack()* prototype for alignment with the  
42109 ISO/IEC 9899:1999 standard.

42110 **NAME**

42111 sigdelset — delete a signal from a signal set

42112 **SYNOPSIS**

42113 CX #include &lt;signal.h&gt;

42114 int sigdelset(sigset\_t \*set, int signo);

42115

42116 **DESCRIPTION**42117 The *sigdelset()* function deletes the individual signal specified by *signo* from the signal set  
42118 pointed to by *set*.42119 Applications should call either *sigemptyset()* or *sigfillset()* at least once for each object of type  
42120 **sigset\_t** prior to any other use of that object. If such an object is not initialized in this way, but is  
42121 nonetheless supplied as an argument to any of *pthread\_sigmask()*, *sigaction()*, *sigaddset()*,  
42122 *sigdelset()*, *sigismember()*, *sigpending()*, *sigprocmask()*, *sigsuspend()*, *sigtimedwait()*, *sigwait()*, or  
42123 *sigwaitinfo()*, the results are undefined.42124 **RETURN VALUE**42125 Upon successful completion, *sigdelset()* shall return 0; otherwise, it shall return -1 and set *errno*  
42126 to indicate the error.42127 **ERRORS**42128 The *sigdelset()* function may fail if:42129 [EINVAL] The *signo* argument is not a valid signal number, or is an unsupported signal  
42130 number.42131 **EXAMPLES**

42132 None.

42133 **APPLICATION USAGE**

42134 None.

42135 **RATIONALE**

42136 None.

42137 **FUTURE DIRECTIONS**

42138 None.

42139 **SEE ALSO**42140 Section 2.4 (on page 478), *sigaction()*, *sigaddset()*, *sigemptyset()*, *sigfillset()*, *sigismember()*,  
42141 *sigpending()*, *sigprocmask()*, *sigsuspend()*, the Base Definitions volume of IEEE Std 1003.1-200x,  
42142 <signal.h>42143 **CHANGE HISTORY**

42144 First released in Issue 3.

42145 Entry included for alignment with the POSIX.1-1988 standard.

42146 **Issue 5**42147 The last paragraph of the DESCRIPTION was included as an APPLICATION USAGE note in  
42148 previous issues.42149 **Issue 6**42150 The SYNOPSIS is marked CX since the presence of this function in the <signal.h> header is an  
42151 extension over the ISO C standard.

42152 **NAME**

42153 sigemptyset — initialize and empty a signal set

42154 **SYNOPSIS**42155 cx `#include <signal.h>`42156 `int sigemptyset(sigset_t *set);`

42157

42158 **DESCRIPTION**42159 The *sigemptyset()* function initializes the signal set pointed to by *set*, such that all signals defined  
42160 in IEEE Std 1003.1-200x are excluded.42161 **RETURN VALUE**42162 Upon successful completion, *sigemptyset()* shall return 0; otherwise, it shall return  $-1$  and set  
42163 *errno* to indicate the error.42164 **ERRORS**

42165 No errors are defined.

42166 **EXAMPLES**

42167 None.

42168 **APPLICATION USAGE**

42169 None.

42170 **RATIONALE**42171 The implementation of the *sigemptyset()* (or *sigfillset()*) function could quite trivially clear (or  
42172 set) all the bits in the signal set. Alternatively, it would be reasonable to initialize part of the  
42173 structure, such as a version field, to permit binary-compatibility between releases where the size  
42174 of the set varies. For such reasons, either *sigemptyset()* or *sigfillset()* must be called prior to any  
42175 other use of the signal set, even if such use is read-only (for example, as an argument to  
42176 *sigpending()*). This function is not intended for dynamic allocation.42177 The *sigfillset()* and *sigemptyset()* functions require that the resulting signal set include (or  
42178 exclude) all the signals defined in this volume of IEEE Std 1003.1-200x. Although it is outside the  
42179 scope of this volume of IEEE Std 1003.1-200x to place this requirement on signals that are  
42180 implemented as extensions, it is recommended that implementation-defined signals also be  
42181 affected by these functions. However, there may be a good reason for a particular signal not to  
42182 be affected. For example, blocking or ignoring an implementation-defined signal may have  
42183 undesirable side effects, whereas the default action for that signal is harmless. In such a case, it  
42184 would be preferable for such a signal to be excluded from the signal set returned by *sigfillset()*.42185 In early proposals there was no distinction between invalid and unsupported signals (the names  
42186 of optional signals that were not supported by an implementation were not defined by that  
42187 implementation). The [EINVAL] error was thus specified as a required error for invalid signals.  
42188 With that distinction, it is not necessary to require implementations of these functions to  
42189 determine whether an optional signal is actually supported, as that could have a significant  
42190 performance impact for little value. The error could have been required for invalid signals and  
42191 optional for unsupported signals, but this seemed unnecessarily complex. Thus, the error is  
42192 optional in both cases.42193 **FUTURE DIRECTIONS**

42194 None.

42195 **SEE ALSO**

42196 Section 2.4 (on page 478), *sigaction()*, *sigaddset()*, *sigdelset()*, *sigfillset()*, *sigismember()*,  
42197 *sigpending()*, *sigprocmask()*, *sigsuspend()*, the Base Definitions volume of IEEE Std 1003.1-200x,  
42198 **<signal.h>**

42199 **CHANGE HISTORY**

42200 First released in Issue 3.

42201 Entry included for alignment with the POSIX.1-1988 standard. |

42202 **Issue 6** |

42203 The SYNOPSIS is marked CX since the presence of this function in the **<signal.h>** header is an |  
42204 extension over the ISO C standard. |

42205 **NAME**

42206 sigfillset — initialize and fill a signal set

42207 **SYNOPSIS**

42208 CX #include &lt;signal.h&gt;

42209 int sigfillset(sigset\_t \*set);

42210

42211 **DESCRIPTION**42212 The *sigfillset()* function shall initialize the signal set pointed to by *set*, such that all signals  
42213 defined in this volume of IEEE Std 1003.1-200x are included.42214 **RETURN VALUE**42215 Upon successful completion, *sigfillset()* shall return 0; otherwise, it shall return -1 and set *errno*  
42216 to indicate the error.42217 **ERRORS**

42218 No errors are defined.

42219 **EXAMPLES**

42220 None.

42221 **APPLICATION USAGE**

42222 None.

42223 **RATIONALE**42224 Refer to *sigemptyset()* (on page 1848).42225 **FUTURE DIRECTIONS**

42226 None.

42227 **SEE ALSO**42228 Section 2.4 (on page 478), *sigaction()*, *sigaddset()*, *sigdelset()*, *sigemptyset()*, *sigismember()*,  
42229 *sigpending()*, *sigprocmask()*, *sigsuspend()*, the Base Definitions volume of IEEE Std 1003.1-200x,  
42230 <signal.h>42231 **CHANGE HISTORY**

42232 First released in Issue 3.

42233 Entry included for alignment with the POSIX.1-1988 standard.

42234 **Issue 6**42235 The SYNOPSIS is marked CX since the presence of this function in the <signal.h> header is an  
42236 extension over the ISO C standard.



42237 **NAME**

42238 sighold, sigignore, sigpause, sigrelse, sigset — signal management

42239 **SYNOPSIS**

```
42240 XSI #include <signal.h>
42241 int sighold(int sig);
42242 int sigignore(int sig);
42243 int sigpause(int sig);
42244 int sigrelse(int sig);
42245 void (*sigset(int sig, void (*disp)(int)))(int);
42246
```

42247 **DESCRIPTION**

42248 Use of any of these functions is unspecified in a multi-threaded process.

42249 The *sighold()*, *sigignore()*, *sigpause()*, *sigrelse()*, and *sigset()* functions provide simplified signal  
42250 management.

42251 The *sigset()* function shall modify signal dispositions. The *sig* argument specifies the signal, |  
42252 which may be any signal except SIGKILL and SIGSTOP. The *disp* argument specifies the signal's |  
42253 disposition, which may be SIG\_DFL, SIG\_IGN, or the address of a signal handler. If *sigset()* is |  
42254 used, and *disp* is the address of a signal handler, the system shall add *sig* to the calling process' |  
42255 signal mask before executing the signal handler; when the signal handler returns, the system |  
42256 shall restore the calling process' signal mask to its state prior to the delivery of the signal. In |  
42257 addition, if *sigset()* is used, and *disp* is equal to SIG\_HOLD, *sig* shall be added to the calling |  
42258 process' signal mask and *sig*'s disposition shall remain unchanged. If *sigset()* is used, and *disp* is |  
42259 not equal to SIG\_HOLD, *sig* shall be removed from the calling process' signal mask.

42260 The *sighold()* function shall add *sig* to the calling process' signal mask. |42261 The *sigrelse()* function shall remove *sig* from the calling process' signal mask. |42262 The *sigignore()* function shall set the disposition of *sig* to SIG\_IGN. |

42263 The *sigpause()* function shall remove *sig* from the calling process' signal mask and suspend the |  
42264 calling process until a signal is received. The *sigpause()* function shall restore the process' signal |  
42265 mask to its original state before returning. |

42266 If the action for the SIGCHLD signal is set to SIG\_IGN, child processes of the calling processes |  
42267 shall not be transformed into zombie processes when they terminate. If the calling process |  
42268 subsequently waits for its children, and the process has no unwaited-for children that were |  
42269 transformed into zombie processes, it shall block until all of its children terminate, and *wait()*, |  
42270 *waitid()*, and *waitpid()* shall fail and set *errno* to [ECHILD].

42271 **RETURN VALUE**

42272 Upon successful completion, *sigset()* shall return SIG\_HOLD if the signal had been blocked and |  
42273 the signal's previous disposition if it had not been blocked. Otherwise, SIG\_ERR shall be |  
42274 returned and *errno* set to indicate the error.

42275 The *sigpause()* function shall suspend execution of the thread until a signal is received, |  
42276 whereupon it shall return -1 and set *errno* to [EINTR].

42277 For all other functions, upon successful completion, 0 shall be returned. Otherwise, -1 shall be |  
42278 returned and *errno* set to indicate the error.

42279 **ERRORS**

42280 These functions shall fail if:

42281 [EINVAL] The *sig* argument is an illegal signal number.

42282 The *sigset()* and *sigignore()* functions shall fail if:

42283 [EINVAL] An attempt is made to catch a signal that cannot be caught, or to ignore a  
42284 signal that cannot be ignored.

42285 **EXAMPLES**

42286 None.

42287 **APPLICATION USAGE**

42288 The *sigaction()* function provides a more comprehensive and reliable mechanism for controlling  
42289 signals; new applications should use *sigaction()* rather than *sigset()*.

42290 The *sighold()* function, in conjunction with *sigrelse()* or *sigpause()*, may be used to establish  
42291 critical regions of code that require the delivery of a signal to be temporarily deferred.

42292 The *sigsuspend()* function should be used in preference to *sigpause()* for broader portability.

42293 **RATIONALE**

42294 None.

42295 **FUTURE DIRECTIONS**

42296 None.

42297 **SEE ALSO**

42298 Section 2.4 (on page 478), *exec*, *pause()*, *sigaction()*, *signal()*, *sigsuspend()*, *waitid()*, the Base  
42299 Definitions volume of IEEE Std 1003.1-200x, <**signal.h**>

42300 **CHANGE HISTORY**

42301 First released in Issue 4 Version 2.

42302 **Issue 5**

42303 Moved from X/OPEN UNIX extension to BASE.

42304 The DESCRIPTION is updated to indicate that the *sigpause()* function restores the process'  
42305 signal mask to its original state before returning.

42306 The RETURN VALUE section is updated to indicate that the *sigpause()* function suspends  
42307 execution of the process until a signal is received, whereupon it returns -1 and sets *errno* to  
42308 [EINTR].

42309 **Issue 6**

42310 The DESCRIPTION is updated to avoid use of the term “must” for application requirements.

42311 References to the *wait3()* function are removed.

42312 The XSI functions are split out into their own reference page.

42313 **NAME**

42314 sigignore — signal management

42315 **SYNOPSIS**42316 XSI `#include <signal.h>`42317 `int sigignore(int sig);`

42318

42319 **DESCRIPTION**42320 Refer to *sighold()*.

## 42321 NAME

42322 siginterrupt — allow signals to interrupt functions

## 42323 SYNOPSIS

42324 XSI #include &lt;signal.h&gt;

42325 int siginterrupt(int sig, int flag);

42326

## 42327 DESCRIPTION

42328 The *siginterrupt()* function shall change the restart behavior when a function is interrupted by |  
42329 the specified signal. The function *siginterrupt(sig, flag)* has an effect as if implemented as:42330 siginterrupt(int sig, int flag) {  
42331 int ret;  
42332 struct sigaction act;  
  
42333 (void) sigaction(sig, NULL, &act);  
42334 if (flag)  
42335 act.sa\_flags &= ~SA\_RESTART;  
42336 else  
42337 act.sa\_flags |= SA\_RESTART;  
42338 ret = sigaction(sig, &act, NULL);  
42339 return ret;  
42340 }

## 42341 RETURN VALUE

42342 Upon successful completion, *siginterrupt()* shall return 0; otherwise, -1 shall be returned and  
42343 *errno* set to indicate the error.

## 42344 ERRORS

42345 The *siginterrupt()* function shall fail if:42346 [EINVAL] The *sig* argument is not a valid signal number.

## 42347 EXAMPLES

42348 None.

## 42349 APPLICATION USAGE

42350 The *siginterrupt()* function supports programs written to historical system interfaces. A |  
42351 conforming application, when being written or rewritten, should use *sigaction()* with the |  
42352 SA\_RESTART flag instead of *siginterrupt()*.

## 42353 RATIONALE

42354 None.

## 42355 FUTURE DIRECTIONS

42356 None.

## 42357 SEE ALSO

42358 Section 2.4 (on page 478), *sigaction()*, the Base Definitions volume of IEEE Std 1003.1-200x,  
42359 <signal.h>

## 42360 CHANGE HISTORY

42361 First released in Issue 4, Version 2.

42362 **Issue 5**

42363 Moved from X/OPEN UNIX extension to BASE.

42364 **NAME**

42365 sigismember — test for a signal in a signal set

42366 **SYNOPSIS**42367 **CX** #include <signal.h>

42368 int sigismember(const sigset\_t \*set, int signo);

42369

42370 **DESCRIPTION**42371 The *sigismember()* function shall test whether the signal specified by *signo* is a member of the set  
42372 pointed to by *set*.42373 Applications should call either *sigemptyset()* or *sigfillset()* at least once for each object of type  
42374 **sigset\_t** prior to any other use of that object. If such an object is not initialized in this way, but is  
42375 nonetheless supplied as an argument to any of *pthread\_sigmask()*, *sigaction()*, *sigaddset()*,  
42376 *sigdelset()*, *sigismember()*, *sigpending()*, *sigprocmask()*, *sigsuspend()*, *sigtimedwait()*, *sigwait()*, or  
42377 *sigwaitinfo()*, the results are undefined.42378 **RETURN VALUE**42379 Upon successful completion, *sigismember()* shall return 1 if the specified signal is a member of  
42380 the specified set, or 0 if it is not. Otherwise, it shall return  $-1$  and set *errno* to indicate the error.42381 **ERRORS**42382 The *sigismember()* function may fail if:42383 [EINVAL] The *signo* argument is not a valid signal number, or is an unsupported signal  
42384 number.42385 **EXAMPLES**

42386 None.

42387 **APPLICATION USAGE**

42388 None.

42389 **RATIONALE**

42390 None.

42391 **FUTURE DIRECTIONS**

42392 None.

42393 **SEE ALSO**42394 Section 2.4 (on page 478), *sigaction()*, *sigaddset()*, *sigdelset()*, *sigfillset()*, *sigemptyset()*,  
42395 *sigpending()*, *sigprocmask()*, *sigsuspend()*, the Base Definitions volume of IEEE Std 1003.1-200x,  
42396 <signal.h>42397 **CHANGE HISTORY**

42398 First released in Issue 3.

42399 Entry included for alignment with the POSIX.1-1988 standard.

42400 **Issue 5**42401 The last paragraph of the DESCRIPTION was included as an APPLICATION USAGE note in  
42402 previous issues.42403 **Issue 6**42404 The SYNOPSIS is marked CX since the presence of this function in the <signal.h> header is an  
42405 extension over the ISO C standard.

42406 **NAME**

42407 siglongjmp — non-local goto with signal handling

42408 **SYNOPSIS**

42409 CX #include &lt;setjmp.h&gt;

42410 void siglongjmp(sigjmp\_buf env, int val);

42411

42412 **DESCRIPTION**42413 The *siglongjmp()* function shall be equivalent to the *longjmp()* function, except as follows:

- 42414 • References to *setjmp()* shall be equivalent to *sigsetjmp()*.
- 42415 • The *siglongjmp()* function shall restore the saved signal mask if and only if the *env* argument
- 42416 was initialized by a call to *sigsetjmp()* with a non-zero *savemask* argument.

42417 **RETURN VALUE**

42418 After *siglongjmp()* is completed, program execution shall continue as if the corresponding  
 42419 invocation of *sigsetjmp()* had just returned the value specified by *val*. The *siglongjmp()* function  
 42420 shall not cause *sigsetjmp()* to return 0; if *val* is 0, *sigsetjmp()* shall return the value 1.

42421 **ERRORS**

42422 No errors are defined.

42423 **EXAMPLES**

42424 None.

42425 **APPLICATION USAGE**

42426 The distinction between *setjmp()* or *longjmp()* and *sigsetjmp()* or *siglongjmp()* is only significant  
 42427 for programs which use *sigaction()*, *sigprocmask()*, or *sigsuspend()*.

42428 **RATIONALE**

42429 None.

42430 **FUTURE DIRECTIONS**

42431 None.

42432 **SEE ALSO**

42433 *longjmp()*, *setjmp()*, *sigprocmask()*, *sigsetjmp()*, *sigsuspend()*, the Base Definitions volume of  
 42434 IEEE Std 1003.1-200x, <setjmp.h>

42435 **CHANGE HISTORY**

42436 First released in Issue 3.

42437 Entry included for alignment with the ISO POSIX-1 standard.

42438 **Issue 5**

42439 The DESCRIPTION is updated for alignment with the POSIX Threads Extension.

42440 **Issue 6**42441 The DESCRIPTION is rewritten in terms of *longjmp()*.

42442 The SYNOPSIS is marked CX since the presence of this function in the <setjmp.h> header is an  
 42443 extension over the ISO C standard.

## 42444 NAME

42445 signal — signal management

## 42446 SYNOPSIS

42447 #include &lt;signal.h&gt;

42448 void (\*signal(int sig, void (\*func)(int)))(int);

## 42449 DESCRIPTION

42450 CX The functionality described on this reference page is aligned with the ISO C standard. Any  
 42451 conflict between the requirements described here and the ISO C standard is unintentional. This  
 42452 volume of IEEE Std 1003.1-200x defers to the ISO C standard.

42453 CX Use of this function is unspecified in a multi-threaded process.

42454 The *signal()* function chooses one of three ways in which receipt of the signal number *sig* is to be  
 42455 subsequently handled. If the value of *func* is SIG\_DFL, default handling for that signal shall  
 42456 occur. If the value of *func* is SIG\_IGN, the signal shall be ignored. Otherwise, the application  
 42457 shall ensure that *func* points to a function to be called when that signal occurs. An invocation of  
 42458 such a function because of a signal, or (recursively) of any further functions called by that  
 42459 invocation (other than functions in the standard library), is called a “signal handler”.

42460 When a signal occurs, and *func* points to a function, it is implementation-defined whether the  
 42461 equivalent of a:

42462 `signal(sig, SIG_DFL);`

42463 is executed or the implementation prevents some implementation-defined set of signals (at least  
 42464 including *sig*) from occurring until the current signal handling has completed. (If the value of *sig*  
 42465 is SIGILL, the implementation may alternatively define that no action is taken.) Next the  
 42466 equivalent of:

42467 `(*func)(sig);`

42468 is executed. If and when the function returns, if the value of *sig* was SIGFPE, SIGILL, or  
 42469 SIGSEGV or any other implementation-defined value corresponding to a computational  
 42470 exception, the behavior is undefined. Otherwise, the program shall resume execution at the  
 42471 CX point it was interrupted. If the signal occurs as the result of calling the *abort()*, *raise()*, *kill()*,  
 42472 *pthread\_kill()*, or *sigqueue()* function, the signal handler shall not call the *raise()* function.

42473 CX If the signal occurs other than as the result of calling *abort()*, *raise()*, *kill()*, *pthread\_kill()*, or  
 42474 *sigqueue()*, the behavior is undefined if the signal handler refers to any object with static storage  
 42475 duration other than by assigning a value to an object declared as volatile **sig\_atomic\_t**, or if the  
 42476 signal handler calls any function in the standard library other than one of the functions listed in  
 42477 Section 2.4 (on page 478). Furthermore, if such a call fails, the value of *errno* is unspecified.

42478 At program start-up, the equivalent of:

42479 `signal(sig, SIG_IGN);`

42480 is executed for some signals, and the equivalent of:

42481 `signal(sig, SIG_DFL);`

42482 CX is executed for all other signals (see *exec*).

## 42483 RETURN VALUE

42484 If the request can be honored, *signal()* shall return the value of *func* for the most recent call to  
 42485 *signal()* for the specified signal *sig*. Otherwise, SIG\_ERR shall be returned and a positive value  
 42486 shall be stored in *errno*.



42487 **ERRORS**42488 The *signal()* function shall fail if:

42489 CX [EINVAL] The *sig* argument is not a valid signal number or an attempt is made to catch a  
 42490 signal that cannot be caught or ignore a signal that cannot be ignored.

42491 The *signal()* function may fail if:

42492 CX [EINVAL] An attempt was made to set the action to SIG\_DFL for a signal that cannot be  
 42493 caught or ignored (or both).

42494 **EXAMPLES**

42495 None.

42496 **APPLICATION USAGE**

42497 The *sigaction()* function provides a more comprehensive and reliable mechanism for controlling  
 42498 signals; new applications should use *sigaction()* rather than *signal()*.

42499 **RATIONALE**

42500 None.

42501 **FUTURE DIRECTIONS**

42502 None.

42503 **SEE ALSO**

42504 Section 2.4 (on page 478), *exec*, *pause()*, *sigaction()*, *sigsuspend()*, *waitid()*, the Base Definitions  
 42505 volume of IEEE Std 1003.1-200x, <**signal.h**>

42506 **CHANGE HISTORY**

42507 First released in Issue 1. Derived from Issue 1 of the SVID.

42508 **Issue 5**

42509 Moved from X/OPEN UNIX extension to BASE.

42510 The DESCRIPTION is updated to indicate that the *sigpause()* function restores the process'  
 42511 signal mask to its original state before returning.

42512 The RETURN VALUE section is updated to indicate that the *sigpause()* function suspends  
 42513 execution of the process until a signal is received, whereupon it returns -1 and sets *errno* to  
 42514 [EINTR].

42515 **Issue 6**

42516 Extensions beyond the ISO C standard are now marked.

42517 The DESCRIPTION is updated to avoid use of the term “must” for application requirements.

42518 The DESCRIPTION is updated for alignment with the ISO/IEC 9899:1999 standard.

42519 References to the *wait3()* function are removed.

42520 The *sighold()*, *sigignore()*, *sigrelse()*, and *sigset()* functions are split out onto their own reference  
 42521 page.

42522 **NAME**

42523 signbit — test sign

42524 **SYNOPSIS**

42525 #include <math.h>

42526 int signbit(real-floating x);

42527 **DESCRIPTION**

42528 cx The functionality described on this reference page is aligned with the ISO C standard. Any  
42529 conflict between the requirements described here and the ISO C standard is unintentional. This  
42530 volume of IEEE Std 1003.1-200x defers to the ISO C standard.

42531 The *signbit()* macro shall determine whether the sign of its argument value is negative. NaNs,  
42532 zeros, and infinities have a sign bit.

42533 **RETURN VALUE**

42534 The *signbit()* macro shall return a non-zero value if and only if the sign of its argument value is  
42535 negative.

42536 **ERRORS**

42537 No errors are defined.

42538 **EXAMPLES**

42539 None.

42540 **APPLICATION USAGE**

42541 None.

42542 **RATIONALE**

42543 None.

42544 **FUTURE DIRECTIONS**

42545 None.

42546 **SEE ALSO**

42547 *fpclassify()*, *isfinite()*, *isinf()*, *isnan()*, *isnormal()*, the Base Definitions volume of  
42548 IEEE Std 1003.1-200x, <math.h>

42549 **CHANGE HISTORY**

42550 First released in Issue 6. Derived from the ISO/IEC 9899:1999 standard.

42551 **NAME**

42552 sigpause — remove a signal from the signal mask and suspend the thread

42553 **SYNOPSIS**

```
42554 xSI #include <signal.h>
```

```
42555 int sigpause(int sig);
```

42556

42557 **DESCRIPTION**

42558 Refer to *sighold()*.

42559 **NAME**

42560 sigpending — examine pending signals

42561 **SYNOPSIS**

42562 CX #include &lt;signal.h&gt;

42563 int sigpending(sigset\_t \*set);

42564

42565 **DESCRIPTION**

42566 The *sigpending()* function shall store, in the location referenced by the *set* argument, the set of  
42567 signals that are blocked from delivery to the calling thread and that are pending on the process  
42568 or the calling thread.

42569 **RETURN VALUE**

42570 Upon successful completion, *sigpending()* shall return 0; otherwise, -1 shall be returned and  
42571 *errno* set to indicate the error.

42572 **ERRORS**

42573 No errors are defined.

42574 **EXAMPLES**

42575 None.

42576 **APPLICATION USAGE**

42577 None.

42578 **RATIONALE**

42579 None.

42580 **FUTURE DIRECTIONS**

42581 None.

42582 **SEE ALSO**

42583 *sigaddset()*, *sigdelset()*, *sigemptyset()*, *sigfillset()*, *sigismember()*, *sigprocmask()*, the Base Definitions  
42584 volume of IEEE Std 1003.1-200x, <signal.h>

42585 **CHANGE HISTORY**

42586 First released in Issue 3.

42587 **Issue 5**

42588 The DESCRIPTION is updated for alignment with the POSIX Threads Extension.

42589 **Issue 6**

42590 The SYNOPSIS is marked CX since the presence of this function in the <signal.h> header is an  
42591 extension over the ISO C standard.

42592 **NAME**

42593 sigprocmask — examine and change blocked signals

42594 **SYNOPSIS**

42595 cx #include &lt;signal.h&gt; |

42596 int sigprocmask(int *how*, const sigset\_t \*restrict *set*,  
42597 sigset\_t \*restrict *oset*);

42598 |

42599 **DESCRIPTION**42600 Refer to *pthread\_sigmask()*.

## 42601 NAME

42602 sigqueue — queue a signal to a process (**REALTIME**)

## 42603 SYNOPSIS

42604 RTS #include &lt;signal.h&gt;

42605 int sigqueue(pid\_t pid, int signo, const union sigval value);

42606

## 42607 DESCRIPTION

42608 The *sigqueue()* function shall cause the signal specified by *signo* to be sent with the value  
 42609 specified by *value* to the process specified by *pid*. If *signo* is zero (the null signal), error checking  
 42610 is performed but no signal is actually sent. The null signal can be used to check the validity of  
 42611 *pid*.

42612 The conditions required for a process to have permission to queue a signal to another process  
 42613 are the same as for the *kill()* function.

42614 The *sigqueue()* function shall return immediately. If SA\_SIGINFO is set for *signo* and if the  
 42615 resources were available to queue the signal, the signal shall be queued and sent to the receiving  
 42616 process. If SA\_SIGINFO is not set for *signo*, then *signo* shall be sent at least once to the receiving  
 42617 process; it is unspecified whether *value* shall be sent to the receiving process as a result of this  
 42618 call.

42619 If the value of *pid* causes *signo* to be generated for the sending process, and if *signo* is not blocked  
 42620 for the calling thread and if no other thread has *signo* unblocked or is waiting in a *sigwait()*  
 42621 function for *signo*, either *signo* or at least the pending, unblocked signal shall be delivered to the  
 42622 calling thread before the *sigqueue()* function returns. Should any multiple pending signals in the  
 42623 range SIGRTMIN to SIGRTMAX be selected for delivery, it shall be the lowest numbered one.  
 42624 The selection order between realtime and non-realtime signals, or between multiple pending  
 42625 non-realtime signals, is unspecified.

## 42626 RETURN VALUE

42627 Upon successful completion, the specified signal shall have been queued, and the *sigqueue()*  
 42628 function shall return a value of zero. Otherwise, the function shall return a value of  $-1$  and set  
 42629 *errno* to indicate the error.

## 42630 ERRORS

42631 The *sigqueue()* function shall fail if:

42632 [EAGAIN] No resources available to queue the signal. The process has already queued  
 42633 SIGQUEUE\_MAX signals that are still pending at the receiver(s), or a system-  
 42634 wide resource limit has been exceeded.

42635 [EINVAL] The value of the *signo* argument is an invalid or unsupported signal number.

42636 [EPERM] The process does not have the appropriate privilege to send the signal to the  
 42637 receiving process.

42638 [ESRCH] The process *pid* does not exist.

42639 **EXAMPLES**

42640 None.

42641 **APPLICATION USAGE**

42642 None.

42643 **RATIONALE**

42644 The *sigqueue()* function allows an application to queue a realtime signal to itself or to another  
42645 process, specifying the application-defined value. This is common practice in realtime  
42646 applications on existing realtime systems. It was felt that specifying another function in the  
42647 *sig...* name space already carved out for signals was preferable to extending the interface to  
42648 *kill()*.

42649 Such a function became necessary when the put/get event function of the message queues was  
42650 removed. It should be noted that the *sigqueue()* function implies reduced performance in a  
42651 security-conscious implementation as the access permissions between the sender and receiver  
42652 have to be checked on each send when the *pid* is resolved into a target process. Such access  
42653 checks were necessary only at message queue open in the previous interface.

42654 The standard developers required that *sigqueue()* have the same semantics with respect to the  
42655 null signal as *kill()*, and that the same permission checking be used. But because of the difficulty  
42656 of implementing the “broadcast” semantic of *kill()* (for example, to process groups) and the  
42657 interaction with resource allocation, this semantic was not adopted. The *sigqueue()* function  
42658 queues a signal to a single process specified by the *pid* argument.

42659 The *sigqueue()* function can fail if the system has insufficient resources to queue the signal. An  
42660 explicit limit on the number of queued signals that a process could send was introduced. While  
42661 the limit is “per-sender”, this volume of IEEE Std 1003.1-200x does not specify that the resources  
42662 be part of the state of the sender. This would require either that the sender be maintained after  
42663 exit until all signals that it had sent to other processes were handled or that all such signals that  
42664 had not yet been acted upon be removed from the queue(s) of the receivers. This volume of  
42665 IEEE Std 1003.1-200x does not preclude this behavior, but an implementation that allocated  
42666 queuing resources from a system-wide pool (with per-sender limits) and that leaves queued  
42667 signals pending after the sender exits is also permitted.

42668 **FUTURE DIRECTIONS**

42669 None.

42670 **SEE ALSO**

42671 Section 2.8.1 (on page 491), the Base Definitions volume of IEEE Std 1003.1-200x, &lt;signal.h&gt;

42672 **CHANGE HISTORY**

42673 First released in Issue 5. Included for alignment with the POSIX Realtime Extension and the  
42674 POSIX Threads Extension.

42675 **Issue 6**42676 The *sigqueue()* function is marked as part of the Realtime Signals Extension option.

42677 The [ENOSYS] error condition has been removed as stubs need not be provided if an  
42678 implementation does not support the Realtime Signals Extension option.

42679 **NAME**

42680 sigrelse — remove a signal from signal mask or modify signal disposition

42681 **SYNOPSIS**

42682 XSI #include <signal.h>

42683 int sigrelse(int sig);

42684

42685 **DESCRIPTION**

42686 Refer to *sighold()*.



42687 **NAME**

42688 sigset — signal management |

42689 **SYNOPSIS**

42690 #include &lt;signal.h&gt;

42691 XSI void (\*sigset(int sig, void (\*disp)(int)))(int);

42692

42693 **DESCRIPTION**42694 Refer to *sighold()*.

42695 **NAME**

42696 sigsetjmp — set jump point for a non-local goto

42697 **SYNOPSIS**42698 `CX #include <setjmp.h>`42699 `int sigsetjmp(sigjmp_buf env, int savemask);`

42700

42701 **DESCRIPTION**42702 The *sigsetjmp()* function shall be equivalent to the *setjmp()* function, except as follows:

- 42703 • References to *setjmp()* are equivalent to *sigsetjmp()*.
- 42704 • References to *longjmp()* are equivalent to *siglongjmp()*.
- 42705 • If the value of the *savemask* argument is not 0, *sigsetjmp()* shall also save the current signal mask of the calling thread as part of the calling environment.

42707 **RETURN VALUE**

42708 If the return is from a successful direct invocation, *sigsetjmp()* shall return 0. If the return is from  
42709 a call to *siglongjmp()*, *sigsetjmp()* shall return a non-zero value.

42710 **ERRORS**

42711 No errors are defined.

42712 **EXAMPLES**

42713 None.

42714 **APPLICATION USAGE**

42715 The distinction between *setjmp()/longjmp()* and *sigsetjmp()/siglongjmp()* is only significant for  
42716 programs which use *sigaction()*, *sigprocmask()*, or *sigsuspend()*.

42717 Note that since this function is defined in terms of *setjmp()*, if *savemask* is zero, it is unspecified  
42718 whether the signal mask is saved.

42719 **RATIONALE**

42720 The ISO C standard specifies various restrictions on the usage of the *setjmp()* macro in order to  
42721 permit implementors to recognize the name in the compiler and not implement an actual  
42722 function. These same restrictions apply to the *sigsetjmp()* macro.

42723 There are processors that cannot easily support these calls, but this was not considered a  
42724 sufficient reason to exclude them.

42725 4.2 BSD, 4.3 BSD, and XSI-conformant systems provide functions named *\_setjmp()* and  
42726 *\_longjmp()* that, together with *setjmp()* and *longjmp()*, provide the same functionality as  
42727 *sigsetjmp()* and *siglongjmp()*. On those systems, *setjmp()* and *longjmp()* save and restore signal  
42728 masks, while *\_setjmp()* and *\_longjmp()* do not. On System V, Release 3 and in corresponding  
42729 issues of the SVID, *setjmp()* and *longjmp()* are explicitly defined not to save and restore signal  
42730 masks. In order to permit existing practice in both cases, the relation of *setjmp()* and *longjmp()* to  
42731 signal masks is not specified, and a new set of functions is defined instead.

42732 The *longjmp()* and *siglongjmp()* functions operate as in the previous issue provided the matching  
42733 *setjmp()* or *sigsetjmp()* has been performed in the same thread. Non-local jumps into contexts  
42734 saved by other threads would be at best a questionable practice and were not considered worthy  
42735 of standardization.

42736 **FUTURE DIRECTIONS**

42737 None.

42738 **SEE ALSO**42739 *siglongjmp()*, *signal()*, *sigprocmask()*, *sigsuspend()*, the Base Definitions volume of  
42740 IEEE Std 1003.1-200x, <**setjmp.h**>42741 **CHANGE HISTORY**

42742 First released in Issue 3.

42743 Entry included for alignment with the POSIX.1-1988 standard.

42744 **Issue 5**

42745 The DESCRIPTION is updated for alignment with the POSIX Threads Extension.

42746 **Issue 6**42747 The DESCRIPTION is reworded in terms of *setjmp()*. |42748 The SYNOPSIS is marked CX since the presence of this function in the <**setjmp.h**> header is an |  
42749 extension over the ISO C standard. |

42750 **NAME**

42751 sigsuspend — wait for a signal

42752 **SYNOPSIS**

42753 cx #include &lt;signal.h&gt;

42754 int sigsuspend(const sigset\_t \*sigmask);

42755

42756 **DESCRIPTION**

42757 The *sigsuspend()* function shall replace the current signal mask of the calling thread with the set  
42758 of signals pointed to by *sigmask* and then suspend the thread until delivery of a signal whose  
42759 action is either to execute a signal-catching function or to terminate the process. This shall not  
42760 cause any other signals that may have been pending on the process to become pending on the  
42761 thread.

42762 If the action is to terminate the process then *sigsuspend()* shall never return. If the action is to  
42763 execute a signal-catching function, then *sigsuspend()* shall return after the signal-catching  
42764 function returns, with the signal mask restored to the set that existed prior to the *sigsuspend()*  
42765 call.

42766 It is not possible to block signals that cannot be ignored. This is enforced by the system without  
42767 causing an error to be indicated.

42768 **RETURN VALUE**

42769 Since *sigsuspend()* suspends thread execution indefinitely, there is no successful completion  
42770 return value. If a return occurs, *-1* shall be returned and *errno* set to indicate the error.

42771 **ERRORS**42772 The *sigsuspend()* function shall fail if:

42773 [EINTR] A signal is caught by the calling process and control is returned from the  
42774 signal-catching function.

42775 **EXAMPLES**

42776 None.

42777 **APPLICATION USAGE**

42778 Normally, at the beginning of a critical code section, a specified set of signals is blocked using  
42779 the *sigprocmask()* function. When the thread has completed the critical section and needs to wait  
42780 for the previously blocked signal(s), it pauses by calling *sigsuspend()* with the mask that was  
42781 returned by the *sigprocmask()* call.

42782 **RATIONALE**

42783 None.

42784 **FUTURE DIRECTIONS**

42785 None.

42786 **SEE ALSO**

42787 Section 2.4 (on page 478), *pause()*, *sigaction()*, *sigaddset()*, *sigdelset()*, *sigemptyset()*, *sigfillset()*, the  
42788 Base Definitions volume of IEEE Std 1003.1-200x, <signal.h>

42789 **CHANGE HISTORY**

42790 First released in Issue 3.

42791 Entry included for alignment with the POSIX.1-1988 standard.

42792 **Issue 5**

42793 The DESCRIPTION is updated for alignment with the POSIX Threads Extension.

42794 **Issue 6**

42795 The text in the RETURN VALUE section has been changed from “suspends process execution”  
42796 to “suspends thread execution”. This reflects IEEE PASC Interpretation 1003.1c #40.

42797 Text in the APPLICATION USAGE section has been replaced. |

42798 The SYNOPSIS is marked CX since the presence of this function in the <signal.h> header is an |  
42799 extension over the ISO C standard. |

## 42800 NAME

42801 sigtimedwait, sigwaitinfo — wait for queued signals (**REALTIME**)

## 42802 SYNOPSIS

42803 RTS #include &lt;signal.h&gt;

```
42804 int sigtimedwait(const sigset_t *restrict set,
42805                 siginfo_t *restrict info,
42806                 const struct timespec *restrict timeout);
42807 int sigwaitinfo(const sigset_t *restrict set,
42808                 siginfo_t *restrict info);
42809
```

## 42810 DESCRIPTION

42811 The *sigtimedwait()* function shall be equivalent to *sigwaitinfo()* except that if none of the signals  
 42812 specified by *set* are pending, *sigtimedwait()* shall wait for the time interval specified in the  
 42813 **timespec** structure referenced by *timeout*. If the **timespec** structure pointed to by *timeout* is  
 42814 zero-valued and if none of the signals specified by *set* are pending, then *sigtimedwait()* shall  
 42815 MON return immediately with an error. If *timeout* is the NULL pointer, the behavior is unspecified. If  
 42816 the Monotonic Clock option is supported, the CLOCK\_MONOTONIC clock shall be used to  
 42817 measure the time interval specified by the *timeout* argument.

42818 The *sigwaitinfo()* function selects the pending signal from the set specified by *set*. Should any of  
 42819 multiple pending signals in the range SIGRTMIN to SIGRTMAX be selected, it shall be the  
 42820 lowest numbered one. The selection order between realtime and non-realtime signals, or  
 42821 between multiple pending non-realtime signals, is unspecified. If no signal in *set* is pending at  
 42822 the time of the call, the calling thread shall be suspended until one or more signals in *set* become  
 42823 pending or until it is interrupted by an unblocked, caught signal.

42824 The *sigwaitinfo()* function shall be equivalent to the *sigwait()* function if the *info* argument is  
 42825 NULL. If the *info* argument is non-NULL, the *sigwaitinfo()* function shall be equivalent to  
 42826 *sigwait()*, except that the selected signal number shall be stored in the *si\_signo* member, and the  
 42827 cause of the signal shall be stored in the *si\_code* member. If any value is queued to the selected  
 42828 signal, the first such queued value shall be dequeued and, if the *info* argument is non-NULL, the  
 42829 value shall be stored in the *si\_value* member of *info*. The system resource used to queue the  
 42830 signal shall be released and returned to the system for other use. If no value is queued, the  
 42831 content of the *si\_value* member is undefined. If no further signals are queued for the selected  
 42832 signal, the pending indication for that signal shall be reset.

## 42833 RETURN VALUE

42834 Upon successful completion (that is, one of the signals specified by *set* is pending or is  
 42835 generated) *sigwaitinfo()* and *sigtimedwait()* shall return the selected signal number. Otherwise,  
 42836 the function shall return a value of -1 and set *errno* to indicate the error.

## 42837 ERRORS

42838 The *sigtimedwait()* function shall fail if:42839 [EAGAIN] No signal specified by *set* was generated within the specified timeout period.42840 The *sigtimedwait()* and *sigwaitinfo()* functions may fail if:

42841 [EINTR] The wait was interrupted by an unblocked, caught signal. It shall be  
 42842 documented in system documentation whether this error causes these  
 42843 functions to fail.

42844 The *sigtimedwait()* function may also fail if:

42845 [EINVAL] The *timeout* argument specified a *tv\_nsec* value less than zero or greater than  
42846 or equal to 1 000 million.

42847 An implementation only checks for this error if no signal is pending in *set* and it is necessary to  
42848 wait.

42849 **EXAMPLES**

42850 None.

42851 **APPLICATION USAGE**

42852 The *sigtimedwait()* function times out and returns an [EAGAIN] error. Application writers  
42853 should note that this is inconsistent with other functions such as *pthread\_cond\_timedwait()* that  
42854 return [ETIMEDOUT].

42855 **RATIONALE**

42856 Existing programming practice on realtime systems uses the ability to pause waiting for a  
42857 selected set of events and handle the first event that occurs in-line instead of in a signal-handling  
42858 function. This allows applications to be written in an event-directed style similar to a state  
42859 machine. This style of programming is useful for largescale transaction processing in which the  
42860 overall throughput of an application and the ability to clearly track states are more important  
42861 than the ability to minimize the response time of individual event handling.

42862 It is possible to construct a signal-waiting macro function out of the realtime signal function  
42863 mechanism defined in this volume of IEEE Std 1003.1-200x. However, such a macro has to  
42864 include the definition of a generalized handler for all signals to be waited on. A significant  
42865 portion of the overhead of handler processing can be avoided if the signal-waiting function is  
42866 provided by the kernel. This volume of IEEE Std 1003.1-200x therefore provides two signal-  
42867 waiting functions—one that waits indefinitely and one with a timeout—as part of the overall  
42868 realtime signal function specification.

42869 The specification of a function with a timeout allows an application to be written that can be  
42870 broken out of a wait after a set period of time if no event has occurred. It was argued that setting  
42871 a timer event before the wait and recognizing the timer event in the wait would also implement  
42872 the same functionality, but at a lower performance level. Because of the performance  
42873 degradation associated with the user-level specification of a timer event and the subsequent  
42874 cancelation of that timer event after the wait completes for a valid event, and the complexity  
42875 associated with handling potential race conditions associated with the user-level method, the  
42876 separate function has been included.

42877 Note that the semantics of the *sigwaitinfo()* function are nearly identical to that of the *sigwait()*  
42878 function defined by this volume of IEEE Std 1003.1-200x. The only difference is that *sigwaitinfo()*  
42879 returns the queued signal value in the *value* argument. The return of the queued value is  
42880 required so that applications can differentiate between multiple events queued to the same  
42881 signal number.

42882 The two distinct functions are being maintained because some implementations may choose to  
42883 implement the POSIX Threads Extension functions and not implement the queued signals  
42884 extensions. Note, though, that *sigwaitinfo()* does not return the queued value if the *value*  
42885 argument is NULL, so the POSIX Threads Extension *sigwait()* function can be implemented as a  
42886 macro on *sigwaitinfo()*.

42887 The *sigtimedwait()* function was separated from the *sigwaitinfo()* function to address concerns  
42888 regarding the overloading of the *timeout* pointer to indicate indefinite wait (no timeout), timed  
42889 wait, and immediate return, and concerns regarding consistency with other functions where the  
42890 conditional and timed waits were separate functions from the pure blocking function. The  
42891 semantics of *sigtimedwait()* are specified such that *sigwaitinfo()* could be implemented as a  
42892 macro with a NULL pointer for *timeout*.

42893 The *sigwait* functions provide a synchronous mechanism for threads to wait for asynchronously  
 42894 generated signals. One important question was how many threads that are suspended in a call to  
 42895 a *sigwait()* function for a signal should return from the call when the signal is sent. Four choices  
 42896 were considered:

- 42897 1. Return an error for multiple simultaneous calls to *sigwait* functions for the same signal.
- 42898 2. One or more threads return.
- 42899 3. All waiting threads return.
- 42900 4. Exactly one thread returns.

42901 Prohibiting multiple calls to *sigwait()* for the same signal was felt to be overly restrictive. The  
 42902 “one or more” behavior made implementation of conforming packages easy at the expense of  
 42903 forcing POSIX threads clients to protect against multiple simultaneous calls to *sigwait()* in  
 42904 application code in order to achieve predictable behavior. There was concern that the “all  
 42905 waiting threads” behavior would result in “signal broadcast storms”, consuming excessive CPU  
 42906 resources by replicating the signals in the general case. Furthermore, no convincing examples  
 42907 could be presented that delivery to all was either simpler or more powerful than delivery to one.

42908 Thus, the consensus was that exactly one thread that was suspended in a call to a *sigwait*  
 42909 function for a signal should return when that signal occurs. This is not an onerous restriction as:

- 42910 • A multi-way signal wait can be built from the single-way wait.
- 42911 • Signals should only be handled by application-level code, as library routines cannot guess  
 42912 what the application wants to do with signals generated for the entire process.
- 42913 • Applications can thus arrange for a single thread to wait for any given signal and call any  
 42914 needed routines upon its arrival.

42915 In an application that is using signals for interprocess communication, signal processing is  
 42916 typically done in one place. Alternatively, if the signal is being caught so that process cleanup  
 42917 can be done, the signal handler thread can call separate process cleanup routines for each  
 42918 portion of the application. Since the application main line started each portion of the application,  
 42919 it is at the right abstraction level to tell each portion of the application to clean up.

42920 Certainly, there exist programming styles where it is logical to consider waiting for a single  
 42921 signal in multiple threads. A simple *sigwait\_multiple()* routine can be constructed to achieve this  
 42922 goal. A possible implementation would be to have each *sigwait\_multiple()* caller registered as  
 42923 having expressed interest in a set of signals. The caller then waits on a thread-specific condition  
 42924 variable. A single server thread calls a *sigwait()* function on the union of all registered signals.  
 42925 When the *sigwait()* function returns, the appropriate state is set and condition variables are  
 42926 broadcast. New *sigwait\_multiple()* callers may cause the pending *sigwait()* call to be canceled  
 42927 and reissued in order to update the set of signals being waited for.

#### 42928 FUTURE DIRECTIONS

42929 None.

#### 42930 SEE ALSO

42931 Section 2.8.1 (on page 491), *pause()*, *pthread\_sigmask()*, *sigaction()*, *sigpending()*, *sigsuspend()*,  
 42932 *sigwait()*, the Base Definitions volume of IEEE Std 1003.1-200x, <**signal.h**>, <**time.h**>

#### 42933 CHANGE HISTORY

42934 First released in Issue 5. Included for alignment with the POSIX Realtime Extension and the  
 42935 POSIX Threads Extension.



42936 **Issue 6**

- 42937 These functions are marked as part of the Realtime Signals Extension option.
- 42938 The Open Group Corrigendum U035/3 is applied. The SYNOPSIS of the *sigwaitinfo()* function  
42939 has been corrected so that the second argument is of type **siginfo\_t\***.
- 42940 The [ENOSYS] error condition has been removed as stubs need not be provided if an  
42941 implementation does not support the Realtime Signals Extension option.
- 42942 The DESCRIPTION is updated for alignment with IEEE Std 1003.1j-2000 by specifying that the  
42943 CLOCK\_MONOTONIC clock, if supported, is used to measure timeout intervals.
- 42944 The **restrict** keyword is added to the *sigtimedwait()* and *sigwaitinfo()* prototypes for alignment  
42945 with the ISO/IEC 9899:1999 standard.

42946 **NAME**

42947 sigwait — wait for queued signals

42948 **SYNOPSIS**42949 `CX` #include <signal.h>42950 `int sigwait(const sigset_t *restrict set, int *restrict sig);`

42951

42952 **DESCRIPTION**

42953 The *sigwait()* function shall select a pending signal from *set*, atomically clear it from the system's  
42954 set of pending signals, and return that signal number in the location referenced by *sig*. If prior to  
42955 the call to *sigwait()* there are multiple pending instances of a single signal number, it is  
42956 implementation-defined whether upon successful return there are any remaining pending  
42957 `RTS` signals for that signal number. If the implementation supports queued signals and there are  
42958 multiple signals queued for the signal number selected, the first such queued signal shall cause a  
42959 return from *sigwait()* and the remainder shall remain queued. If no signal in *set* is pending at the  
42960 time of the call, the thread shall be suspended until one or more becomes pending. The signals  
42961 defined by *set* shall have been blocked at the time of the call to *sigwait()*; otherwise, the behavior  
42962 is undefined. The effect of *sigwait()* on the signal actions for the signals in *set* is unspecified.

42963 If more than one thread is using *sigwait()* to wait for the same signal, no more than one of these  
42964 threads shall return from *sigwait()* with the signal number. Which thread returns from *sigwait()*  
42965 if more than a single thread is waiting is unspecified.

42966 `RTS` Should any of the multiple pending signals in the range SIGRTMIN to SIGRTMAX be selected, it  
42967 shall be the lowest numbered one. The selection order between realtime and non-realtime  
42968 signals, or between multiple pending non-realtime signals, is unspecified.

42969 **RETURN VALUE**

42970 Upon successful completion, *sigwait()* shall store the signal number of the received signal at the  
42971 location referenced by *sig* and return zero. Otherwise, an error number shall be returned to  
42972 indicate the error.

42973 **ERRORS**42974 The *sigwait()* function may fail if:42975 [EINVAL] The *set* argument contains an invalid or unsupported signal number.42976 **EXAMPLES**

42977 None.

42978 **APPLICATION USAGE**

42979 None.

42980 **RATIONALE**

42981 To provide a convenient way for a thread to wait for a signal, this volume of  
42982 IEEE Std 1003.1-200x provides the *sigwait()* function. For most cases where a thread has to wait  
42983 for a signal, the *sigwait()* function should be quite convenient, efficient, and adequate.

42984 However, requests were made for a lower-level primitive than *sigwait()* and for semaphores that  
42985 could be used by threads. After some consideration, threads were allowed to use semaphores  
42986 and *sem\_post()* was defined to be async-signal and async-cancel-safe.

42987 In summary, when it is necessary for code run in response to an asynchronous signal to notify a  
42988 thread, *sigwait()* should be used to handle the signal. Alternatively, if the implementation  
42989 provides semaphores, they also can be used, either following *sigwait()* or from within a signal  
42990 handling routine previously registered with *sigaction()*.

42991 **FUTURE DIRECTIONS**

42992 None.

42993 **SEE ALSO**

42994 Section 2.4 (on page 478), Section 2.8.1 (on page 491), *pause()*, *pthread\_sigmask()*, *sigaction()*,  
42995 *sigpending()*, *sigsuspend()*, *sigwaitinfo()*, the Base Definitions volume of IEEE Std 1003.1-200x,  
42996 <**signal.h**>, <**time.h**>

42997 **CHANGE HISTORY**

42998 First released in Issue 5. Included for alignment with the POSIX Realtime Extension and the  
42999 POSIX Threads Extension.

43000 **Issue 6**

43001 The RATIONALE section is added.

43002 The **restrict** keyword is added to the *sigwait()* prototype for alignment with the  
43003 ISO/IEC 9899:1999 standard.

43004 **NAME**

43005 sigwaitinfo — wait for queued signals (**REALTIME**)

43006 **SYNOPSIS**

43007 RTS `#include <signal.h>`

43008 `int sigwaitinfo(const sigset_t *restrict set, siginfo_t *restrict info);`

43009

43010 **DESCRIPTION**

43011 Refer to *sigtimedwait()*.

43012 **NAME**

43013 sin, sinf, sinl — sine function

43014 **SYNOPSIS**

43015 #include &lt;math.h&gt;

43016 double sin(double x);

43017 float sinf(float x);

43018 long double sinl(long double x);

43019 **DESCRIPTION**

43020 **CX** The functionality described on this reference page is aligned with the ISO C standard. Any  
 43021 conflict between the requirements described here and the ISO C standard is unintentional. This  
 43022 volume of IEEE Std 1003.1-200x defers to the ISO C standard.

43023 These functions shall compute the sine of their argument *x*, measured in radians.

43024 An application wishing to check for error situations should set *errno* to zero and call  
 43025 *feclearexcept*(FE\_ALL\_EXCEPT) before calling these functions. On return, if *errno* is non-zero or  
 43026 *fetestexcept*(FE\_INVALID | FE\_DIVBYZERO | FE\_OVERFLOW | FE\_UNDERFLOW) is non-  
 43027 zero, an error has occurred.

43028 **RETURN VALUE**43029 Upon successful completion, these functions shall return the sine of *x*.43030 **MX** If *x* is NaN, a NaN shall be returned.43031 If *x* is  $\pm 0$ , *x* shall be returned.43032 If *x* is subnormal, a range error may occur and *x* should be returned.

43033 If *x* is  $\pm\text{Inf}$ , a domain error shall occur, and either a NaN (if supported), or an implementation-  
 43034 defined value shall be returned.

43035 **ERRORS**

43036 These functions shall fail if:

43037 **MX** **Domain Error** The *x* argument is  $\pm\text{Inf}$ .

43038 If the integer expression (math\_errhandling & MATH\_ERRNO) is non-zero, |  
 43039 then *errno* shall be set to [EDOM]. If the integer expression (math\_errhandling |  
 43040 & MATH\_ERREXCEPT) is non-zero, then the invalid floating-point exception |  
 43041 shall be raised. |

43042 These functions may fail if:

43043 **MX** **Range Error** The value of *x* is subnormal

43044 If the integer expression (math\_errhandling & MATH\_ERRNO) is non-zero, |  
 43045 then *errno* shall be set to [ERANGE]. If the integer expression |  
 43046 (math\_errhandling & MATH\_ERREXCEPT) is non-zero, then the underflow |  
 43047 floating-point exception shall be raised. |

43048 **EXAMPLES**43049 **Taking the Sine of a 45-Degree Angle**

```
43050 #include <math.h>
43051 ...
43052 double radians = 45.0 * M_PI / 180;
43053 double result;
43054 ...
43055 result = sin(radians);
```

43056 **APPLICATION USAGE**

43057 These functions may lose accuracy when their argument is near a multiple of  $\pi$  or is far from 0.0.

43058 On error, the expressions (math\_errhandling & MATH\_ERRNO) and (math\_errhandling &  
43059 MATH\_ERREXCEPT) are independent of each other, but at least one of them must be non-zero.

43060 **RATIONALE**

43061 None.

43062 **FUTURE DIRECTIONS**

43063 None.

43064 **SEE ALSO**

43065 *asin()*, *feclearexcept()*, *fetestexcept()*, *isnan()*, the Base Definitions volume of IEEE Std 1003.1-200x, |  
43066 Section 4.18, Treatment of Error Conditions for Mathematical Functions, <math.h> |

43067 **CHANGE HISTORY**

43068 First released in Issue 1. Derived from Issue 1 of the SVID.

43069 **Issue 5**

43070 The last two paragraphs of the DESCRIPTION were included as APPLICATION USAGE notes  
43071 in previous issues.

43072 **Issue 6**

43073 The *sinf()* and *sinl()* functions are added for alignment with the ISO/IEC 9899:1999 standard.

43074 The DESCRIPTION, RETURN VALUE, ERRORS, and APPLICATION USAGE sections are  
43075 revised to align with the ISO/IEC 9899:1999 standard.

43076 IEC 60559:1989 standard floating-point extensions over the ISO/IEC 9899:1999 standard are  
43077 marked.

43078 **NAME**

43079       sinf — sine function

43080 **SYNOPSIS**

43081       #include <math.h>

43082       float sinf(float x);

43083 **DESCRIPTION**

43084       Refer to *sin()*.

43085 **NAME**

43086       sinh, sinhf, sinhl — hyperbolic sine function

43087 **SYNOPSIS**

43088       #include &lt;math.h&gt;

43089       double sinh(double x);

43090       float sinhf(float x);

43091       long double sinhl(long double x);

43092 **DESCRIPTION**43093 **CX**       The functionality described on this reference page is aligned with the ISO C standard. Any  
43094 conflict between the requirements described here and the ISO C standard is unintentional. This  
43095 volume of IEEE Std 1003.1-200x defers to the ISO C standard.43096       These functions shall compute the hyperbolic sine of their argument *x*.43097       An application wishing to check for error situations should set *errno* to zero and call  
43098 *feclearexcept*(FE\_ALL\_EXCEPT) before calling these functions. On return, if *errno* is non-zero or  
43099 *fetestexcept*(FE\_INVALID | FE\_DIVBYZERO | FE\_OVERFLOW | FE\_UNDERFLOW) is non-  
43100 zero, an error has occurred.43101 **RETURN VALUE**43102       Upon successful completion, these functions shall return the hyperbolic sine of *x*.43103       If the result would cause an overflow, a range error shall occur and  $\pm$ HUGE\_VAL,  
43104  $\pm$ HUGE\_VALF, and  $\pm$ HUGE\_VALL (with the same sign as *x*) shall be returned as appropriate for  
43105 the type of the function.43106 **MX**       If *x* is NaN, a NaN shall be returned.43107       If *x* is  $\pm 0$ , or  $\pm$ Inf, *x* shall be returned.43108       If *x* is subnormal, a range error may occur and *x* should be returned.43109 **ERRORS**

43110       These functions shall fail if:

43111       Range Error       The result would cause an overflow.

43112               If the integer expression (math\_errhandling & MATH\_ERRNO) is non-zero, |  
43113 then *errno* shall be set to [ERANGE]. If the integer expression |  
43114 (math\_errhandling & MATH\_ERREXCEPT) is non-zero, then the overflow |  
43115 floating-point exception shall be raised. |

43116       These functions may fail if:

43117 **MX**       Range Error       The value *x* is subnormal.43118               If the integer expression (math\_errhandling & MATH\_ERRNO) is non-zero, |  
43119 then *errno* shall be set to [ERANGE]. If the integer expression |  
43120 (math\_errhandling & MATH\_ERREXCEPT) is non-zero, then the underflow |  
43121 floating-point exception shall be raised. |



43122 **EXAMPLES**

43123 None.

43124 **APPLICATION USAGE**

43125 On error, the expressions (math\_errhandling & MATH\_ERRNO) and (math\_errhandling &  
43126 MATH\_ERREXCEPT) are independent of each other, but at least one of them must be non-zero.

43127 **RATIONALE**

43128 None.

43129 **FUTURE DIRECTIONS**

43130 None.

43131 **SEE ALSO**

43132 *asinh()*, *cosh()*, *feclearexcept()*, *fetestexcept()*, *isnan()*, *tanh()*, the Base Definitions volume of |  
43133 IEEE Std 1003.1-200x, Section 4.18, Treatment of Error Conditions for Mathematical Functions, |  
43134 <math.h>

43135 **CHANGE HISTORY**

43136 First released in Issue 1. Derived from Issue 1 of the SVID.

43137 **Issue 5**

43138 The DESCRIPTION is updated to indicate how an application should check for an error. This  
43139 text was previously published in the APPLICATION USAGE section.

43140 **Issue 6**43141 The *sinhf()* and *sinhl()* functions are added for alignment with the ISO/IEC 9899:1999 standard.

43142 The DESCRIPTION, RETURN VALUE, ERRORS, and APPLICATION USAGE sections are  
43143 revised to align with the ISO/IEC 9899:1999 standard.

43144 IEC 60559:1989 standard floating-point extensions over the ISO/IEC 9899:1999 standard are  
43145 marked.

43146 **NAME**

43147       sinl — sine function

43148 **SYNOPSIS**

43149       #include <math.h>

43150       long double sinl(long double x);

43151 **DESCRIPTION**

43152       Refer to *sin()*.

43153 **NAME**

43154 sleep — suspend execution for an interval of time

43155 **SYNOPSIS**

43156 #include &lt;unistd.h&gt;

43157 unsigned sleep(unsigned *seconds*);43158 **DESCRIPTION**

43159 The *sleep()* function shall cause the calling thread to be suspended from execution until either  
 43160 the number of realtime seconds specified by the argument *seconds* has elapsed or a signal is  
 43161 delivered to the calling thread and its action is to invoke a signal-catching function or to  
 43162 terminate the process. The suspension time may be longer than requested due to the scheduling  
 43163 of other activity by the system.

43164 If a SIGALRM signal is generated for the calling process during execution of *sleep()* and if the  
 43165 SIGALRM signal is being ignored or blocked from delivery, it is unspecified whether *sleep()*  
 43166 returns when the SIGALRM signal is scheduled. If the signal is being blocked, it is also  
 43167 unspecified whether it remains pending after *sleep()* returns or it is discarded.

43168 If a SIGALRM signal is generated for the calling process during execution of *sleep()*, except as a  
 43169 result of a prior call to *alarm()*, and if the SIGALRM signal is not being ignored or blocked from  
 43170 delivery, it is unspecified whether that signal has any effect other than causing *sleep()* to return.

43171 If a signal-catching function interrupts *sleep()* and examines or changes either the time a  
 43172 SIGALRM is scheduled to be generated, the action associated with the SIGALRM signal, or  
 43173 whether the SIGALRM signal is blocked from delivery, the results are unspecified.

43174 If a signal-catching function interrupts *sleep()* and calls *siglongjmp()* or *longjmp()* to restore an  
 43175 environment saved prior to the *sleep()* call, the action associated with the SIGALRM signal and  
 43176 the time at which a SIGALRM signal is scheduled to be generated are unspecified. It is also  
 43177 unspecified whether the SIGALRM signal is blocked, unless the process' signal mask is restored  
 43178 as part of the environment.

43179 XSI Interactions between *sleep()* and any of *setitimer()*, *ualarm()*, or *usleep()* are unspecified.

43180 **RETURN VALUE**

43181 If *sleep()* returns because the requested time has elapsed, the value returned shall be 0. If *sleep()*  
 43182 returns due to delivery of a signal, the return value shall be the “unslept” amount (the requested  
 43183 time minus the time actually slept) in seconds.

43184 **ERRORS**

43185 No errors are defined.

43186 **EXAMPLES**

43187 None.

43188 **APPLICATION USAGE**

43189 None.

43190 **RATIONALE**

43191 There are two general approaches to the implementation of the *sleep()* function. One is to use the  
 43192 *alarm()* function to schedule a SIGALRM signal and then suspend the process waiting for that  
 43193 signal. The other is to implement an independent facility. This volume of IEEE Std 1003.1-200x  
 43194 permits either approach.

43195 In order to comply with the requirement that no primitive shall change a process attribute unless  
 43196 explicitly described by this volume of IEEE Std 1003.1-200x, an implementation using SIGALRM  
 43197 must carefully take into account any SIGALRM signal scheduled by previous *alarm()* calls, the

43198 action previously established for SIGALRM, and whether SIGALRM was blocked. If a SIGALRM  
43199 has been scheduled before the *sleep()* would ordinarily complete, the *sleep()* must be shortened  
43200 to that time and a SIGALRM generated (possibly simulated by direct invocation of the signal-  
43201 catching function) before *sleep()* returns. If a SIGALRM has been scheduled after the *sleep()*  
43202 would ordinarily complete, it must be rescheduled for the same time before *sleep()* returns. The  
43203 action and blocking for SIGALRM must be saved and restored.

43204 Historical implementations often implement the SIGALRM-based version using *alarm()* and  
43205 *pause()*. One such implementation is prone to infinite hangups, as described in *pause()*. Another  
43206 such implementation uses the C-language *setjmp()* and *longjmp()* functions to avoid that  
43207 window. That implementation introduces a different problem: when the SIGALRM signal  
43208 interrupts a signal-catching function installed by the user to catch a different signal, the  
43209 *longjmp()* aborts that signal-catching function. An implementation based on *sigprocmask()*,  
43210 *alarm()*, and *sigsuspend()* can avoid these problems.

43211 Despite all reasonable care, there are several very subtle, but detectable and unavoidable,  
43212 differences between the two types of implementations. These are the cases mentioned in this  
43213 volume of IEEE Std 1003.1-200x where some other activity relating to SIGALRM takes place, and  
43214 the results are stated to be unspecified. All of these cases are sufficiently unusual as not to be of  
43215 concern to most applications.

43216 See also the discussion of the term *realtime* in *alarm()*.

43217 Since *sleep()* can be implemented using *alarm()*, the discussion about alarms occurring early  
43218 under *alarm()* applies to *sleep()* as well.

43219 Application writers should note that the type of the argument *seconds* and the return value of  
43220 *sleep()* is **unsigned**. That means that a Strictly Conforming POSIX System Interfaces Application  
43221 cannot pass a value greater than the minimum guaranteed value for {UINT\_MAX}, which the  
43222 ISO C standard sets as 65 535, and any application passing a larger value is restricting its  
43223 portability. A different type was considered, but historical implementations, including those  
43224 with a 16-bit **int** type, consistently use either **unsigned** or **int**.

43225 Scheduling delays may cause the process to return from the *sleep()* function significantly after  
43226 the requested time. In such cases, the return value should be set to zero, since the formula  
43227 (requested time minus the time actually spent) yields a negative number and *sleep()* returns an  
43228 **unsigned**.

#### 43229 FUTURE DIRECTIONS

43230 None.

#### 43231 SEE ALSO

43232 *alarm()*, *getitimer()*, *nanosleep()*, *pause()*, *sigaction()*, *sigsetjmp()*, *ualarm()*, *usleep()*, the Base  
43233 Definitions volume of IEEE Std 1003.1-200x, <**unistd.h**>

#### 43234 CHANGE HISTORY

43235 First released in Issue 1. Derived from Issue 1 of the SVID.

#### 43236 Issue 5

43237 The DESCRIPTION is updated for alignment with the POSIX Threads Extension.

43238 **NAME**43239            *snprintf* — print formatted output43240 **SYNOPSIS**

43241            #include &lt;stdio.h&gt;

43242            int snprintf(char \*restrict *s*, size\_t *n*,43243                const char \*restrict *format*, ...);43244 **DESCRIPTION**43245            Refer to *fprintf*().

43246 **NAME**

43247 socket — create an endpoint for communication

43248 **SYNOPSIS**

43249 #include &lt;sys/socket.h&gt;

43250 int socket(int *domain*, int *type*, int *protocol*);43251 **DESCRIPTION**43252 The *socket()* function shall create an unbound socket in a communications domain, and return a  
43253 file descriptor that can be used in later function calls that operate on sockets.43254 The *socket()* function takes the following arguments:43255 *domain* Specifies the communications domain in which a socket is to be created.43256 *type* Specifies the type of socket to be created.43257 *protocol* Specifies a particular protocol to be used with the socket. Specifying a *protocol*  
43258 of 0 causes *socket()* to use an unspecified default protocol appropriate for the  
43259 requested socket type.43260 The *domain* argument specifies the address family used in the communications domain. The  
43261 address families supported by the system are implementation-defined.43262 Symbolic constants that can be used for the domain argument are defined in the <sys/socket.h>  
43263 header.43264 The *type* argument specifies the socket type, which determines the semantics of communication  
43265 over the socket. The following socket types are defined; implementations may specify additional  
43266 socket types:43267 SOCK\_STREAM Provides sequenced, reliable, bidirectional, connection-mode byte  
43268 streams, and may provide a transmission mechanism for out-of-band  
43269 data.43270 SOCK\_DGRAM Provides datagrams, which are connectionless-mode, unreliable messages  
43271 of fixed maximum length.43272 SOCK\_SEQPACKET Provides sequenced, reliable, bidirectional, connection-mode  
43273 transmission path for records. A record can be sent using one or more  
43274 output operations and received using one or more input operations, but a  
43275 single operation never transfers part of more than one record. Record  
43276 boundaries are visible to the receiver via the MSG\_EOR flag.43277 If the *protocol* argument is non-zero, it shall specify a protocol that is supported by the address  
43278 family. If the *protocol* argument is zero, the default protocol for this address family and type shall  
43279 be used. The protocols supported by the system are implementation-defined.43280 The process may need to have appropriate privileges to use the *socket()* function or to create  
43281 some sockets.43282 **RETURN VALUE**43283 Upon successful completion, *socket()* shall return a non-negative integer, the socket file  
43284 descriptor. Otherwise, a value of -1 shall be returned and *errno* set to indicate the error.43285 **ERRORS**43286 The *socket()* function shall fail if:

43287 [EAFNOSUPPORT]

43288 The implementation does not support the specified address family.

- 43289 [EMFILE] No more file descriptors are available for this process.
- 43290 [ENFILE] No more file descriptors are available for the system.
- 43291 [EPROTONOSUPPORT]  
 43292 The protocol is not supported by the address family, or the protocol is not  
 43293 supported by the implementation.
- 43294 [EPROTOTYPE] The socket type is not supported by the protocol.
- 43295 The *socket()* function may fail if:
- 43296 [EACCES] The process does not have appropriate privileges.
- 43297 [ENOBUFS] Insufficient resources were available in the system to perform the operation.
- 43298 [ENOMEM] Insufficient memory was available to fulfill the request.
- 43299 **EXAMPLES**
- 43300 None.
- 43301 **APPLICATION USAGE**
- 43302 The documentation for specific address families specifies which protocols each address family  
 43303 supports. The documentation for specific protocols specifies which socket types each protocol  
 43304 supports.
- 43305 The application can determine whether an address family is supported by trying to create a  
 43306 socket with *domain* set to the protocol in question.
- 43307 **RATIONALE**
- 43308 None.
- 43309 **FUTURE DIRECTIONS**
- 43310 None.
- 43311 **SEE ALSO**
- 43312 *accept()*, *bind()*, *connect()*, *getsockname()*, *getsockopt()*, *listen()*, *recv()*, *recvfrom()*, *recvmsg()*,  
 43313 *send()*, *sendmsg()*, *setsockopt()*, *shutdown()*, *socketpair()*, the Base Definitions volume of  
 43314 IEEE Std 1003.1-200x, <[netinet/in.h](#)>, <[sys/socket.h](#)>
- 43315 **CHANGE HISTORY**
- 43316 First released in Issue 6. Derived from the XNS, Issue 5.2 specification.

43317 **NAME**

43318 socketpair — create a pair of connected sockets

43319 **SYNOPSIS**

43320 #include &lt;sys/socket.h&gt;

43321 int socketpair(int *domain*, int *type*, int *protocol*,  
43322 int *socket\_vector*[2]);43323 **DESCRIPTION**43324 The *socketpair()* function shall create an unbound pair of connected sockets in a specified *domain*,  
43325 of a specified *type*, under the protocol optionally specified by the *protocol* argument. The two  
43326 sockets shall be identical. The file descriptors used in referencing the created sockets shall be  
43327 returned in *socket\_vector*[0] and *socket\_vector*[1].43328 The *socketpair()* function takes the following arguments:

43329	<i>domain</i>	Specifies the communications domain in which the sockets are to be created.
43330	<i>type</i>	Specifies the type of sockets to be created.
43331	<i>protocol</i>	Specifies a particular protocol to be used with the sockets. Specifying a
43332		<i>protocol</i> of 0 causes <i>socketpair()</i> to use an unspecified default protocol
43333		appropriate for the requested socket type.
43334	<i>socket_vector</i>	Specifies a 2-integer array to hold the file descriptors of the created socket
43335		pair.

43336 The *type* argument specifies the socket type, which determines the semantics of communications  
43337 over the socket. The following socket types are defined; implementations may specify additional  
43338 socket types:

43339	SOCK_STREAM	Provides sequenced, reliable, bidirectional, connection-mode byte
43340		streams, and may provide a transmission mechanism for out-of-band
43341		data.
43342	SOCK_DGRAM	Provides datagrams, which are connectionless-mode, unreliable messages
43343		of fixed maximum length.
43344	SOCK_SEQPACKET	Provides sequenced, reliable, bidirectional, connection-mode
43345		transmission paths for records. A record can be sent using one or more
43346		output operations and received using one or more input operations, but a
43347		single operation never transfers part of more than one record. Record
43348		boundaries are visible to the receiver via the MSG_EOR flag.

43349 If the *protocol* argument is non-zero, it shall specify a protocol that is supported by the address  
43350 family. If the *protocol* argument is zero, the default protocol for this address family and type shall  
43351 be used. The protocols supported by the system are implementation-defined.43352 The process may need to have appropriate privileges to use the *socketpair()* function or to create  
43353 some sockets.43354 **RETURN VALUE**43355 Upon successful completion, this function shall return 0; otherwise, -1 shall be returned and  
43356 *errno* set to indicate the error.43357 **ERRORS**43358 The *socketpair()* function shall fail if:

43359 [EAFNOSUPPORT]

43360 The implementation does not support the specified address family.



- 43361 [EMFILE] No more file descriptors are available for this process.
- 43362 [ENFILE] No more file descriptors are available for the system.
- 43363 [EOPNOTSUPP] The specified protocol does not permit creation of socket pairs.
- 43364 [EPROTONOSUPPORT]  
43365 The protocol is not supported by the address family, or the protocol is not  
43366 supported by the implementation.
- 43367 [EPROTOTYPE] The socket type is not supported by the protocol.
- 43368 The *socketpair()* function may fail if:
- 43369 [EACCES] The process does not have appropriate privileges.
- 43370 [ENOBUFS] Insufficient resources were available in the system to perform the operation.
- 43371 [ENOMEM] Insufficient memory was available to fulfill the request.

**43372 EXAMPLES**

43373 None.

**43374 APPLICATION USAGE**

43375 The documentation for specific address families specifies which protocols each address family  
43376 supports. The documentation for specific protocols specifies which socket types each protocol  
43377 supports.

43378 The *socketpair()* function is used primarily with UNIX domain sockets and need not be  
43379 supported for other domains.

**43380 RATIONALE**

43381 None.

**43382 FUTURE DIRECTIONS**

43383 None.

**43384 SEE ALSO**

43385 *socket()*, the Base Definitions volume of IEEE Std 1003.1-200x, <sys/socket.h>

**43386 CHANGE HISTORY**

43387 First released in Issue 6. Derived from the XNS, Issue 5.2 specification.

43388 **NAME**

43389       printf — print formatted output

43390 **SYNOPSIS**

43391       #include <stdio.h>

43392       int printf(char \*restrict *s*, const char \*restrict *format*, ...);

43393 **DESCRIPTION**

43394       Refer to *fprintf()*.

43395 **NAME**

43396 sqrt, sqrtf, sqrtl — square root function

43397 **SYNOPSIS**

43398 #include &lt;math.h&gt;

43399 double sqrt(double x);

43400 float sqrtf(float x);

43401 long double sqrtl(long double x);

43402 **DESCRIPTION**

43403 **CX** The functionality described on this reference page is aligned with the ISO C standard. Any  
 43404 conflict between the requirements described here and the ISO C standard is unintentional. This  
 43405 volume of IEEE Std 1003.1-200x defers to the ISO C standard.

43406 These functions shall compute the square root of their argument  $x$ ,  $\sqrt{x}$ .

43407 An application wishing to check for error situations should set *errno* to zero and call  
 43408 *feclearexcept*(FE\_ALL\_EXCEPT) before calling these functions. On return, if *errno* is non-zero or  
 43409 *fetestexcept*(FE\_INVALID | FE\_DIVBYZERO | FE\_OVERFLOW | FE\_UNDERFLOW) is non-  
 43410 zero, an error has occurred.

43411 **RETURN VALUE**43412 Upon successful completion, these functions shall return the square root of  $x$ .

43413 **MX** For finite values of  $x < -0$ , a domain error shall occur, and either a NaN (if supported), or an  
 43414 implementation-defined value shall be returned.

43415 **MX** If  $x$  is NaN, a NaN shall be returned.

43416 If  $x$  is  $\pm 0$ , or  $+\text{Inf}$ ,  $x$  shall be returned.

43417 If  $x$  is  $-\text{Inf}$ , a domain error shall occur, and either a NaN (if supported), or an implementation-  
 43418 defined value shall be returned.

43419 **ERRORS**

43420 These functions shall fail if:

43421 **MX** Domain Error The finite value of  $x$  is  $< -0$ , or  $x$  is  $-\text{Inf}$ .

43422 If the integer expression (math\_errhandling & MATH\_ERRNO) is non-zero, |  
 43423 then *errno* shall be set to [EDOM]. If the integer expression (math\_errhandling |  
 43424 & MATH\_ERREXCEPT) is non-zero, then the invalid floating-point exception |  
 43425 shall be raised. |

43426 **EXAMPLES**43427 **Taking the Square Root of 9.0**

43428 #include &lt;math.h&gt;

43429 ...

43430 double x = 9.0;

43431 double result;

43432 ...

43433 result = sqrt(x);

43434 **APPLICATION USAGE**

43435           On error, the expressions (math\_errhandling & MATH\_ERRNO) and (math\_errhandling &  
43436           MATH\_ERREXCEPT) are independent of each other, but at least one of them must be non-zero.

43437 **RATIONALE**

43438           None.

43439 **FUTURE DIRECTIONS**

43440           None.

43441 **SEE ALSO**

43442           *feclearexcept()*, *fetetestexcept()*, *isnan()*, the Base Definitions volume of IEEE Std 1003.1-200x, |  
43443           Section 4.18, Treatment of Error Conditions for Mathematical Functions, <math.h>, <stdio.h> |

43444 **CHANGE HISTORY**

43445           First released in Issue 1. Derived from Issue 1 of the SVID.

43446 **Issue 5**

43447           The DESCRIPTION is updated to indicate how an application should check for an error. This  
43448           text was previously published in the APPLICATION USAGE section.

43449 **Issue 6**

43450           The *sqrtrf()* and *sqrtrl()* functions are added for alignment with the ISO/IEC 9899:1999 standard.

43451           The DESCRIPTION, RETURN VALUE, ERRORS, and APPLICATION USAGE sections are  
43452           revised to align with the ISO/IEC 9899:1999 standard.

43453           IEC 60559:1989 standard floating-point extensions over the ISO/IEC 9899:1999 standard are  
43454           marked.

43455 **NAME**

43456           srand — pseudo-random number generator

43457 **SYNOPSIS**

43458           #include <stdlib.h>

43459           void srand(unsigned *seed*);

43460 **DESCRIPTION**

43461           Refer to *rand()*.

43462 **NAME**

43463       srand48 — seed uniformly distributed double-precision pseudo-random number generator

43464 **SYNOPSIS**

43465 XSI       #include &lt;stdlib.h&gt;

43466       void srand48(long *seedval*);

43467

43468 **DESCRIPTION**43469       Refer to *drand48()*.

43470 **NAME**

43471           srandom — seed pseudo-random number generator

43472 **SYNOPSIS**

43473 XSI       #include &lt;stdlib.h&gt;

43474           void srandom(unsigned *seed*);

43475

43476 **DESCRIPTION**43477           Refer to *initstate()*.

43478 **NAME**43479        `sscanf` — convert formatted input43480 **SYNOPSIS**43481        `#include <stdio.h>`43482        `int sscanf(const char *restrict s, const char *restrict format, ...);`43483 **DESCRIPTION**43484        Refer to *fscanf()*.



43485 **NAME**

43486       stat — get file status

43487 **SYNOPSIS**

43488       #include &lt;sys/stat.h&gt;

43489       int stat(const char \*restrict path, struct stat \*restrict buf);

43490 **DESCRIPTION**

43491       The *stat()* function shall obtain information about the named file and write it to the area pointed  
 43492       to by the *buf* argument. The *path* argument points to a pathname naming a file. Read, write, or  
 43493       execute permission of the named file is not required. An implementation that provides  
 43494       additional or alternate file access control mechanisms may, under implementation-defined  
 43495       conditions, cause *stat()* to fail. In particular, the system may deny the existence of the file  
 43496       specified by *path*.

43497       If the named file is a symbolic link, the *stat()* function shall continue pathname resolution using  
 43498       the contents of the symbolic link, and shall return information pertaining to the resulting file if  
 43499       the file exists.

43500       The *buf* argument is a pointer to a **stat** structure, as defined in the <sys/stat.h> header, into  
 43501       which information is placed concerning the file.

43502       The *stat()* function shall update any time-related fields (as described in the Base Definitions  
 43503       volume of IEEE Std 1003.1-200x, Section 4.7, File Times Update), before writing into the **stat**  
 43504       structure.

43505       The structure members *st\_mode*, *st\_ino*, *st\_dev*, *st\_uid*, *st\_gid*, *st\_atime*, *st\_ctime*, and *st\_mtime*  
 43506       shall have meaningful values for all file types defined in this volume of IEEE Std 1003.1-200x.  
 43507       The value of the member *st\_nlink* shall be set to the number of links to the file.

43508 **RETURN VALUE**

43509       Upon successful completion, 0 shall be returned. Otherwise, -1 shall be returned and *errno* set to  
 43510       indicate the error.

43511 **ERRORS**43512       The *stat()* function shall fail if:

43513       [EACCES]       Search permission is denied for a component of the path prefix.

43514       [EIO]         An error occurred while reading from the file system.

43515       [ELOOP]       A loop exists in symbolic links encountered during resolution of the *path*  
 43516       argument.

43517       [ENAMETOOLONG]

43518       The length of the *path* argument exceeds {PATH\_MAX} or a pathname  
 43519       component is longer than {NAME\_MAX}.43520       [ENOENT]       A component of *path* does not name an existing file or *path* is an empty string.

43521       [ENOTDIR]      A component of the path prefix is not a directory.

43522       [EOVERFLOW]   The file size in bytes or the number of blocks allocated to the file or the file  
 43523       serial number cannot be represented correctly in the structure pointed to by  
 43524       *buf*.43525       The *stat()* function may fail if:43526       [ELOOP]       More than {SYMLOOP\_MAX} symbolic links were encountered during  
 43527       resolution of the *path* argument.

43528 [ENAMETOOLONG]  
 43529 As a result of encountering a symbolic link in resolution of the *path* argument, |  
 43530 the length of the substituted pathname string exceeded {PATH\_MAX}. |

43531 [EOVERFLOW] A value to be stored would overflow one of the members of the **stat** structure.

43532 **EXAMPLES**43533 **Obtaining File Status Information**

43534 The following example shows how to obtain file status information for a file named  
 43535 **/home/cnd/mod1**. The structure variable *buffer* is defined for the **stat** structure.

```
43536 #include <sys/types.h>
43537 #include <sys/stat.h>
43538 #include <fcntl.h>

43539 struct stat buffer;
43540 int status;
43541 ...
43542 status = stat("/home/cnd/mod1", &buffer);
```

43543 **Getting Directory Information**

43544 The following example fragment gets status information for each entry in a directory. The call to  
 43545 the *stat()* function stores file information in the **stat** structure pointed to by *statbuf*. The lines  
 43546 that follow the *stat()* call format the fields in the **stat** structure for presentation to the user of the  
 43547 program.

```
43548 #include <sys/types.h>
43549 #include <sys/stat.h>
43550 #include <dirent.h>
43551 #include <pwd.h>
43552 #include <grp.h>
43553 #include <time.h>
43554 #include <locale.h>
43555 #include <langinfo.h>
43556 #include <stdio.h>
43557 #include <stdint.h>

43558 struct dirent *dp;
43559 struct stat statbuf;
43560 struct passwd *pwd;
43561 struct group *grp;
43562 struct tm *tm;
43563 char datestring[256];
43564 ...
43565 /* Loop through directory entries */
43566 while ((dp = readdir(dir)) != NULL) {
43567     /* Get entry's information. */
43568     if (stat(dp->d_name, &statbuf) == -1)
43569         continue;
43570     /* Print out type, permissions, and number of links. */
43571     printf("%10.10s", sparm (statbuf.st_mode));
43572     printf("%4d", statbuf.st_nlink);
```

```

43573     /* Print out owners name if it is found using getpwuid(). */
43574     if ((pwd = getpwuid(statbuf.st_uid)) != NULL)
43575         printf(" %-8.8s", pwd->pw_name);
43576     else
43577         printf(" %-8d", statbuf.st_uid);
43578     /* Print out group name if it's found using getgrgid(). */
43579     if ((grp = getgrgid(statbuf.st_gid)) != NULL)
43580         printf(" %-8.8s", grp->gr_name);
43581     else
43582         printf(" %-8d", statbuf.st_gid);
43583     /* Print size of file. */
43584     printf(" %9jd", (intmax_t)statbuf.st_size);
43585     tm = localtime(&statbuf.st_mtime);
43586     /* Get localized date string. */
43587     strftime(datestring, sizeof(datestring), nl_langinfo(D_T_FMT), tm);
43588     printf(" %s %s\n", datestring, dp->d_name);
43589 }

```

#### 43590 APPLICATION USAGE

43591 None.

#### 43592 RATIONALE

43593 The intent of the paragraph describing “additional or alternate file access control mechanisms”  
 43594 is to allow a secure implementation where a process with a label that does not dominate the  
 43595 file’s label cannot perform a *stat()* function. This is not related to read permission; a process with  
 43596 a label that dominates the file’s label does not need read permission. An implementation that  
 43597 supports write-up operations could fail *lstat()* function calls even though it has a valid file  
 43598 descriptor open for writing.

#### 43599 FUTURE DIRECTIONS

43600 None.

#### 43601 SEE ALSO

43602 *lstat()*, *lstat()*, *readlink()*, *symlink()*, the Base Definitions volume of IEEE Std 1003.1-200x,  
 43603 <sys/stat.h>, <sys/types.h>

#### 43604 CHANGE HISTORY

43605 First released in Issue 1. Derived from Issue 1 of the SVID.

#### 43606 Issue 5

43607 Large File Summit extensions are added.

#### 43608 Issue 6

43609 In the SYNOPSIS, the optional include of the <sys/types.h> header is removed.

43610 The following new requirements on POSIX implementations derive from alignment with the  
 43611 Single UNIX Specification:

- 43612 • The requirement to include <sys/types.h> has been removed. Although <sys/types.h> was  
 43613 required for conforming implementations of previous POSIX specifications, it was not  
 43614 required for UNIX applications.
- 43615 • The [EIO] mandatory error condition is added.

- 43616 • The [ELOOP] mandatory error condition is added.
  - 43617 • The [EOVERFLOW] mandatory error condition is added. This change is to support large
  - 43618 files.
  - 43619 • The [ENAMETOOLONG] and the second [EOVERFLOW] optional error conditions are
  - 43620 added.
- 43621 The following changes were made to align with the IEEE P1003.1a draft standard:
- 43622 • Details are added regarding the treatment of symbolic links.
  - 43623 • The [ELOOP] optional error condition is added.
- 43624 The DESCRIPTION is updated to avoid use of the term “must” for application requirements.
- 43625 The **restrict** keyword is added to the *stat()* prototype for alignment with the ISO/IEC 9899:1999
- 43626 standard.

43627 **NAME**

43628           statvfs — get file system information

43629 **SYNOPSIS**

43630 xSI        #include &lt;sys/statvfs.h&gt;

43631            int statvfs(const char \*restrict *path*, struct statvfs \*restrict *buf*);

43632

43633 **DESCRIPTION**43634            Refer to *fstatvfs()*.

43635 **NAME**

43636 stderr, stdin, stdout — standard I/O streams

43637 **SYNOPSIS**

43638 #include &lt;stdio.h&gt;

43639 extern FILE \*stderr, \*stdin, \*stdout;

43640 **DESCRIPTION**

43641 cx The functionality described on this reference page is aligned with the ISO C standard. Any  
43642 conflict between the requirements described here and the ISO C standard is unintentional. This  
43643 volume of IEEE Std 1003.1-200x defers to the ISO C standard.

43644 A file with associated buffering is called a *stream* and is declared to be a pointer to a defined type  
43645 **FILE**. The *fopen()* function shall create certain descriptive data for a stream and return a pointer  
43646 to designate the stream in all further transactions. Normally, there are three open streams with  
43647 constant pointers declared in the <stdio.h> header and associated with the standard open files.

43648 At program start-up, three streams shall be predefined and need not be opened explicitly:  
43649 *standard input* (for reading conventional input), *standard output* (for writing conventional output),  
43650 and *standard error* (for writing diagnostic output). When opened, the standard error stream is not  
43651 fully buffered; the standard input and standard output streams are fully buffered if and only if  
43652 the stream can be determined not to refer to an interactive device.

43653 cx The following symbolic values in <unistd.h> define the file descriptors that shall be associated  
43654 with the C-language *stdin*, *stdout*, and *stderr* when the application is started:

43655 STDIN\_FILENO Standard input value, *stdin*. Its value is 0.

43656 STDOUT\_FILENO Standard output value, *stdout*. Its value is 1.

43657 STDERR\_FILENO Standard error value, *stderr*. Its value is 2.

43658 The *stderr* stream is expected to be open for reading and writing.

43659 **RETURN VALUE**

43660 None.

43661 **ERRORS**

43662 No errors are defined.

43663 **EXAMPLES**

43664 None.

43665 **APPLICATION USAGE**

43666 None.

43667 **RATIONALE**

43668 None.

43669 **FUTURE DIRECTIONS**

43670 None.

43671 **SEE ALSO**

43672 *fclose()*, *feof()*, *ferror()*, *fileno()*, *fopen()*, *fread()*, *fseek()*, *getc()*, *gets()*, *popen()*, *printf()*, *putc()*,  
43673 *puts()*, *read()*, *scanf()*, *setbuf()*, *setvbuf()*, *tmpfile()*, *ungetc()*, *vprintf()*, the Base Definitions  
43674 volume of IEEE Std 1003.1-200x, <stdio.h>, <unistd.h>

43675 **CHANGE HISTORY**

43676 First released in Issue 1.

43677 **Issue 6**

43678 Extensions beyond the ISO C standard are now marked.

43679 A note that *stderr* is expected to be open for reading and writing is added to the DESCRIPTION.

43680 **NAME**

43681 strcasecmp, strncasecmp — case-insensitive string comparisons

43682 **SYNOPSIS**43683 XSI `#include <strings.h>`43684 `int strcasecmp(const char *s1, const char *s2);`43685 `int strncasecmp(const char *s1, const char *s2, size_t n);`

43686

43687 **DESCRIPTION**

43688 The *strcasecmp()* function shall compare, while ignoring differences in case, the string pointed to  
43689 by *s1* to the string pointed to by *s2*. The *strncasecmp()* function shall compare, while ignoring  
43690 differences in case, not more than *n* bytes from the string pointed to by *s1* to the string pointed to  
43691 by *s2*.

43692 In the POSIX locale, *strcasecmp()* and *strncasecmp()* shall behave as if the strings had been |  
43693 converted to lowercase and then a byte comparison performed. The results are unspecified in |  
43694 other locales. |

43695 **RETURN VALUE**

43696 Upon completion, *strcasecmp()* shall return an integer greater than, equal to, or less than 0, if the  
43697 string pointed to by *s1* is, ignoring case, greater than, equal to, or less than the string pointed to  
43698 by *s2*, respectively.

43699 Upon successful completion, *strncasecmp()* shall return an integer greater than, equal to, or less  
43700 than 0, if the possibly null-terminated array pointed to by *s1* is, ignoring case, greater than, equal  
43701 to, or less than the possibly null-terminated array pointed to by *s2*, respectively.

43702 **ERRORS**

43703 No errors are defined.

43704 **EXAMPLES**

43705 None.

43706 **APPLICATION USAGE**

43707 None.

43708 **RATIONALE**

43709 None.

43710 **FUTURE DIRECTIONS**

43711 None.

43712 **SEE ALSO**

43713 The Base Definitions volume of IEEE Std 1003.1-200x, &lt;strings.h&gt;

43714 **CHANGE HISTORY**

43715 First released in Issue 4, Version 2.

43716 **Issue 5**

43717 Moved from X/OPEN UNIX extension to BASE.



43718 **NAME**

43719           strcat — concatenate two strings

43720 **SYNOPSIS**

43721           #include &lt;string.h&gt;

43722           char \*strcat(char \*restrict *s1*, const char \*restrict *s2*);43723 **DESCRIPTION**

43724 **CX**       The functionality described on this reference page is aligned with the ISO C standard. Any  
43725       conflict between the requirements described here and the ISO C standard is unintentional. This  
43726       volume of IEEE Std 1003.1-200x defers to the ISO C standard.

43727       The *strcat()* function shall append a copy of the string pointed to by *s2* (including the  
43728       terminating null byte) to the end of the string pointed to by *s1*. The initial byte of *s2* overwrites  
43729       the null byte at the end of *s1*. If copying takes place between objects that overlap, the behavior is  
43730       undefined.

43731 **RETURN VALUE**43732       The *strcat()* function shall return *s1*; no return value is reserved to indicate an error.43733 **ERRORS**

43734       No errors are defined.

43735 **EXAMPLES**

43736       None.

43737 **APPLICATION USAGE**

43738       This issue is aligned with the ISO C standard; this does not affect compatibility with XPG3  
43739       applications. Reliable error detection by this function was never guaranteed.

43740 **RATIONALE**

43741       None.

43742 **FUTURE DIRECTIONS**

43743       None.

43744 **SEE ALSO**43745       *strncat()*, the Base Definitions volume of IEEE Std 1003.1-200x, <string.h>43746 **CHANGE HISTORY**

43747       First released in Issue 1. Derived from Issue 1 of the SVID.

43748 **Issue 6**43749       The *strcat()* prototype is updated for alignment with the ISO/IEC 9899:1999 standard.

43750 **NAME**

43751           strchr — string scanning operation

43752 **SYNOPSIS**

43753           #include &lt;string.h&gt;

43754           char \*strchr(const char \*s, int c);

43755 **DESCRIPTION**43756 **CX**       The functionality described on this reference page is aligned with the ISO C standard. Any  
43757       conflict between the requirements described here and the ISO C standard is unintentional. This  
43758       volume of IEEE Std 1003.1-200x defers to the ISO C standard.43759 **CX**       The *strchr()* function shall locate the first occurrence of *c* (converted to an **unsigned char**) in the  
43760       string pointed to by *s*. The terminating null byte is considered to be part of the string.43761 **RETURN VALUE**43762           Upon completion, *strchr()* shall return a pointer to the byte, or a null pointer if the byte was not  
43763           found.43764 **ERRORS**

43765           No errors are defined.

43766 **EXAMPLES**

43767           None.

43768 **APPLICATION USAGE**

43769           None.

43770 **RATIONALE**

43771           None.

43772 **FUTURE DIRECTIONS**

43773           None.

43774 **SEE ALSO**43775           *strchr()*, the Base Definitions volume of IEEE Std 1003.1-200x, <string.h>43776 **CHANGE HISTORY**

43777           First released in Issue 1. Derived from Issue 1 of the SVID.

43778 **Issue 6**

43779           Extensions beyond the ISO C standard are now marked.

43780 **NAME**

43781 strcmp — compare two strings

43782 **SYNOPSIS**

43783 #include &lt;string.h&gt;

43784 int strcmp(const char \*s1, const char \*s2);

43785 **DESCRIPTION**

43786 cx The functionality described on this reference page is aligned with the ISO C standard. Any  
 43787 conflict between the requirements described here and the ISO C standard is unintentional. This  
 43788 volume of IEEE Std 1003.1-200x defers to the ISO C standard.

43789 The *strcmp()* function shall compare the string pointed to by *s1* to the string pointed to by *s2*.

43790 The sign of a non-zero return value shall be determined by the sign of the difference between the  
 43791 values of the first pair of bytes (both interpreted as type **unsigned char**) that differ in the strings  
 43792 being compared.

43793 **RETURN VALUE**

43794 Upon completion, *strcmp()* shall return an integer greater than, equal to, or less than 0, if the  
 43795 string pointed to by *s1* is greater than, equal to, or less than the string pointed to by *s2*,  
 43796 respectively.

43797 **ERRORS**

43798 No errors are defined.

43799 **EXAMPLES**43800 **Checking a Password Entry**

43801 The following example compares the information read from standard input to the value of the  
 43802 name of the user entry. If the *strcmp()* function returns 0 (indicating a match), a further check  
 43803 will be made to see if the user entered the proper old password. The *crypt()* function shall  
 43804 encrypt the old password entered by the user, using the value of the encrypted password in the  
 43805 **passwd** structure as the salt. If this value matches the value of the encrypted **passwd** in the  
 43806 structure, the entered password *oldpasswd* is the correct user's password. Finally, the program  
 43807 encrypts the new password so that it can store the information in the **passwd** structure.

```

43808 #include <string.h>
43809 #include <unistd.h>
43810 #include <stdio.h>
43811 ...
43812 int valid_change;
43813 struct passwd *p;
43814 char user[100];
43815 char oldpasswd[100];
43816 char newpasswd[100];
43817 char savepasswd[100];
43818 ...
43819 if (strcmp(p->pw_name, user) == 0) {
43820     if (strcmp(p->pw_passwd, crypt(oldpasswd, p->pw_passwd)) == 0) {
43821         strcpy(savepasswd, crypt(newpasswd, user));
43822         p->pw_passwd = savepasswd;
43823         valid_change = 1;
43824     }
43825     else {
```

```
43826             fprintf(stderr, "Old password is not valid\n");
43827         }
43828     }
43829     ...
```

43830 **APPLICATION USAGE**

43831 None.

43832 **RATIONALE**

43833 None.

43834 **FUTURE DIRECTIONS**

43835 None.

43836 **SEE ALSO**

43837 *strncmp()*, the Base Definitions volume of IEEE Std 1003.1-200x, <**string.h**>

43838 **CHANGE HISTORY**

43839 First released in Issue 1. Derived from Issue 1 of the SVID.

43840 **Issue 6**

43841 Extensions beyond the ISO C standard are now marked.

43842 **NAME**

43843 strcoll — string comparison using collating information

43844 **SYNOPSIS**

43845 #include &lt;string.h&gt;

43846 int strcoll(const char \*s1, const char \*s2);

43847 **DESCRIPTION**

43848 CX The functionality described on this reference page is aligned with the ISO C standard. Any  
 43849 conflict between the requirements described here and the ISO C standard is unintentional. This  
 43850 volume of IEEE Std 1003.1-200x defers to the ISO C standard.

43851 The *strcoll()* function shall compare the string pointed to by *s1* to the string pointed to by *s2*,  
 43852 both interpreted as appropriate to the *LC\_COLLATE* category of the current locale.

43853 CX The *strcoll()* function shall not change the setting of *errno* if successful.

43854 Since no return value is reserved to indicate an error, an application wishing to check for error  
 43855 situations should set *errno* to 0, then call *strcoll()*, then check *errno*.

43856 **RETURN VALUE**

43857 Upon successful completion, *strcoll()* shall return an integer greater than, equal to, or less than 0,  
 43858 according to whether the string pointed to by *s1* is greater than, equal to, or less than the string  
 43859 CX pointed to by *s2* when both are interpreted as appropriate to the current locale. On error,  
 43860 *strcoll()* may set *errno*, but no return value is reserved to indicate an error.

43861 **ERRORS**43862 The *strcoll()* function may fail if:

43863 CX [EINVAL] The *s1* or *s2* arguments contain characters outside the domain of the collating  
 43864 sequence.

43865 **EXAMPLES**43866 **Comparing Nodes**

43867 The following example uses an application-defined function, *node\_compare()*, to compare two  
 43868 nodes based on an alphabetical ordering of the *string* field.

```
43869 #include <string.h>
43870 ...
43871 struct node { /* These are stored in the table. */
43872     char *string;
43873     int length;
43874 };
43875 ...
43876 int node_compare(const void *node1, const void *node2)
43877 {
43878     return strcoll(((const struct node *)node1)->string,
43879                 ((const struct node *)node2)->string);
43880 }
43881 ...
```

43882 **APPLICATION USAGE**43883 The *strxfrm()* and *strcmp()* functions should be used for sorting large lists.

43884 **RATIONALE**

43885           None.

43886 **FUTURE DIRECTIONS**

43887           None.

43888 **SEE ALSO**43889           *strcmp()*, *strxfrm()*, the Base Definitions volume of IEEE Std 1003.1-200x, <**string.h**>43890 **CHANGE HISTORY**

43891           First released in Issue 3.

43892 **Issue 5**43893           The DESCRIPTION is updated to indicate that *errno* does not be changed if the function is  
43894           successful.43895 **Issue 6**

43896           Extensions beyond the ISO C standard are now marked.

43897           The following new requirements on POSIX implementations derive from alignment with the  
43898           Single UNIX Specification:

- 43899
- The [EINVAL] optional error condition is added.

43900           An example is added. |

43901 **NAME**

43902           strcpy — copy a string

43903 **SYNOPSIS**

43904           #include &lt;string.h&gt;

43905           char \*strcpy(char \*restrict *s1*, const char \*restrict *s2*);43906 **DESCRIPTION**

43907 *cx*       The functionality described on this reference page is aligned with the ISO C standard. Any  
 43908 conflict between the requirements described here and the ISO C standard is unintentional. This  
 43909 volume of IEEE Std 1003.1-200x defers to the ISO C standard.

43910       The *strcpy()* function shall copy the string pointed to by *s2* (including the terminating null byte)  
 43911 into the array pointed to by *s1*. If copying takes place between objects that overlap, the behavior  
 43912 is undefined.

43913 **RETURN VALUE**43914       The *strcpy()* function shall return *s1*; no return value is reserved to indicate an error.43915 **ERRORS**

43916       No errors are defined.

43917 **EXAMPLES**43918           **Initializing a String**43919       The following example copies the string "-----" into the *permstring* variable.

```
43920       #include <string.h>
43921       ...
43922       static char permstring[11];
43923       ...
43924       strcpy(permstring, "-----");
43925       ...
```

43926           **Storing a Key and Data**

43927       The following example allocates space for a key using *malloc()* then uses *strcpy()* to place the  
 43928 key there. Then it allocates space for data using *malloc()*, and uses *strcpy()* to place data there.  
 43929 (The user-defined function *dbfree()* frees memory previously allocated to an array of type **struct**  
 43930 **element** \*.)

```
43931       #include <string.h>
43932       #include <stdlib.h>
43933       #include <stdio.h>
43934       ...
43935       /* Structure used to read data and store it. */
43936       struct element {
43937           char *key;
43938           char *data;
43939       };
43940       struct element *tbl, *curtbl;
43941       char *key, *data;
43942       int count;
43943       ...
43944       void dbfree(struct element *, int);
```

```
43945     ...
43946     if ((curtbl->key = malloc(strlen(key) + 1)) == NULL) {
43947         perror("malloc"); dbfree(tbl, count); return NULL;
43948     }
43949     strcpy(curtbl->key, key);
43950
43951     if ((curtbl->data = malloc(strlen(data) + 1)) == NULL) {
43952         perror("malloc"); free(curtbl->key); dbfree(tbl, count); return NULL;
43953     }
43954     strcpy(curtbl->data, data);
43955     ...
```

**43955 APPLICATION USAGE**

43956 Character movement is performed differently in different implementations. Thus, overlapping  
43957 moves may yield surprises.

43958 This issue is aligned with the ISO C standard; this does not affect compatibility with XPG3  
43959 applications. Reliable error detection by this function was never guaranteed.

**43960 RATIONALE**

43961 None.

**43962 FUTURE DIRECTIONS**

43963 None.

**43964 SEE ALSO**

43965 *strncpy()*, the Base Definitions volume of IEEE Std 1003.1-200x, <**string.h**>

**43966 CHANGE HISTORY**

43967 First released in Issue 1. Derived from Issue 1 of the SVID.

**43968 Issue 6**

43969 The *strcpy()* prototype is updated for alignment with the ISO/IEC 9899:1999 standard.



43970 **NAME**

43971 strcspn — get length of a complementary substring

43972 **SYNOPSIS**

43973 #include &lt;string.h&gt;

43974 size\_t strcspn(const char \*s1, const char \*s2);

43975 **DESCRIPTION**

43976 cx The functionality described on this reference page is aligned with the ISO C standard. Any  
 43977 conflict between the requirements described here and the ISO C standard is unintentional. This  
 43978 volume of IEEE Std 1003.1-200x defers to the ISO C standard.

43979 The *strcspn()* function shall compute the length (in bytes) of the maximum initial segment of the  
 43980 string pointed to by *s1* which consists entirely of bytes *not* from the string pointed to by *s2*.

43981 **RETURN VALUE**

43982 The *strcspn()* function shall return the length of the computed segment of the string pointed to  
 43983 by *s1*; no return value is reserved to indicate an error.

43984 **ERRORS**

43985 No errors are defined.

43986 **EXAMPLES**

43987 None.

43988 **APPLICATION USAGE**

43989 None.

43990 **RATIONALE**

43991 None.

43992 **FUTURE DIRECTIONS**

43993 None.

43994 **SEE ALSO**43995 *strspn()*, the Base Definitions volume of IEEE Std 1003.1-200x, <**string.h**>43996 **CHANGE HISTORY**

43997 First released in Issue 1. Derived from Issue 1 of the SVID.

43998 **Issue 5**

43999 The RETURN VALUE section is updated to indicated that *strcspn()* returns the length of *s1*, and  
 44000 not *s1* itself as was previously stated.

44001 **Issue 6**

44002 The Open Group Corrigendum U030/1 is applied. The text of the RETURN VALUE section is  
 44003 updated to indicate that the computed segment length is returned, not the *s1* length.

44004 **NAME**

44005        **strdup** — duplicate a string

44006 **SYNOPSIS**

44007 XSI        #include <string.h>

44008        char \*strdup(const char \*s1);

44009

44010 **DESCRIPTION**

44011        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.

44014 **RETURN VALUE**

44015        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.

44017 **ERRORS**

44018        The *strdup()* function may fail if:

44019        [ENOMEM]       Storage space available is insufficient.

44020 **EXAMPLES**

44021        None.

44022 **APPLICATION USAGE**

44023        None.

44024 **RATIONALE**

44025        None.

44026 **FUTURE DIRECTIONS**

44027        None.

44028 **SEE ALSO**

44029        *free()*, *malloc()*, the Base Definitions volume of IEEE Std 1003.1-200x, <string.h>

44030 **CHANGE HISTORY**

44031        First released in Issue 4, Version 2.

44032 **Issue 5**

44033        Moved from X/OPEN UNIX extension to BASE.

44034 **NAME**

44035            strerror, strerror\_r — get error message string

44036 **SYNOPSIS**

44037            #include &lt;string.h&gt;

44038            char \*strerror(int *errnum*);44039 TSF        int strerror\_r(int *errnum*, char \**strerrbuf*, size\_t *buflen*);

44040

44041 **DESCRIPTION**44042 CX        For *strerror()*: 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.44045            The *strerror()* function shall map the error number in *errnum* to a locale-dependent error message string and shall return a pointer to it. Typically, the values for *errnum* come from *errno*, but *strerror()* shall map any value of type **int** to a message.44048            The string pointed to shall not be modified by the application, but may be overwritten by a subsequent call to *strerror()* or *perror()*.44049 CX        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.44050 CX        The implementation shall behave as if no function defined in this volume of IEEE Std 1003.1-200x calls *strerror()*.44051            The *strerror()* function shall not change the setting of *errno* if successful.44052            Since 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*.44053            The *strerror()* function need not be reentrant. A function that is not required to be reentrant is not required to be thread-safe.44054 CX        The *strerror\_r()* function shall map the error number in *errnum* to a locale-dependent error message string and shall return the string in the buffer pointed to by *strerrbuf*, with length *buflen*.

44055

44056 TSF        The *strerror\_r()* function shall map the error number in *errnum* to a locale-dependent error message string and shall return the string in the buffer pointed to by *strerrbuf*, with length *buflen*.44057            Upon successful completion, *strerror\_r()* shall return 0. Otherwise, an error number shall be returned to indicate the error.44058            Upon successful completion, *strerror\_r()* shall return 0. Otherwise, an error number shall be returned to indicate the error.44059 TSF        The *strerror\_r()* function shall map the error number in *errnum* to a locale-dependent error message string and shall return the string in the buffer pointed to by *strerrbuf*, with length *buflen*.44060            The *strerror\_r()* function shall map the error number in *errnum* to a locale-dependent error message string and shall return the string in the buffer pointed to by *strerrbuf*, with length *buflen*.44061            The *strerror\_r()* function shall map the error number in *errnum* to a locale-dependent error message string and shall return the string in the buffer pointed to by *strerrbuf*, with length *buflen*.44062            The *strerror\_r()* function shall map the error number in *errnum* to a locale-dependent error message string and shall return the string in the buffer pointed to by *strerrbuf*, with length *buflen*.44063            The *strerror\_r()* function shall map the error number in *errnum* to a locale-dependent error message string and shall return the string in the buffer pointed to by *strerrbuf*, with length *buflen*.44064            The *strerror\_r()* function shall map the error number in *errnum* to a locale-dependent error message string and shall return the string in the buffer pointed to by *strerrbuf*, with length *buflen*.

44065

44066            The *strerror\_r()* function shall map the error number in *errnum* to a locale-dependent error message string and shall return the string in the buffer pointed to by *strerrbuf*, with length *buflen*.44067            The *strerror\_r()* function shall map the error number in *errnum* to a locale-dependent error message string and shall return the string in the buffer pointed to by *strerrbuf*, with length *buflen*.

44068

44073 **EXAMPLES**

44074           None.

44075 **APPLICATION USAGE**

44076           None.

44077 **RATIONALE**

44078           None.

44079 **FUTURE DIRECTIONS**

44080           None.

44081 **SEE ALSO**44082           *perror()*, the Base Definitions volume of IEEE Std 1003.1-200x, <**string.h**>44083 **CHANGE HISTORY**

44084           First released in Issue 3.

44085 **Issue 5**44086           The DESCRIPTION is updated to indicate that *errno* is not changed if the function is successful.

44087           A note indicating that this function need not be reentrant is added to the DESCRIPTION.

44088 **Issue 6**

44089           Extensions beyond the ISO C standard are now marked.

44090           The following new requirements on POSIX implementations derive from alignment with the  
44091           Single UNIX Specification:

- 44092
- In the RETURN VALUE section, the fact that *errno* may be set is added.
  - The [EINVAL] optional error condition is added.

44094           The DESCRIPTION is updated to avoid use of the term “must” for application requirements.

44095           The *strerror\_r()* function is added in response to IEEE PASC Interpretation 1003.1c #39.44096           The *strerror\_r()* function is marked as part of the Thread-Safe Functions option.

44097 **NAME**

44098 strfmon — convert monetary value to a string

44099 **SYNOPSIS**

44100 xSI #include &lt;monetary.h&gt;

44101 ssize\_t strfmon(char \*restrict *s*, size\_t *maxsize*,  
44102 const char \*restrict *format*, ...);

44103

44104 **DESCRIPTION**44105 The *strfmon()* function shall place characters into the array pointed to by *s* as controlled by the  
44106 string pointed to by *format*. No more than *maxsize* bytes are placed into the array.44107 The format is a character string, beginning and ending in its initial state, if any, that contains two |  
44108 types of objects: *plain characters*, which are simply copied to the output stream, and *conversion* |  
44109 *specifications*, each of which shall result in the fetching of zero or more arguments which are |  
44110 converted and formatted. The results are undefined if there are insufficient arguments for the |  
44111 format. If the format is exhausted while arguments remain, the excess arguments are simply |  
44112 ignored.

44113 The application shall ensure that a conversion specification consists of the following sequence:

- 44114 • A '%' character
- 
- 44115 • Optional flags
- 
- 44116 • Optional field width
- 
- 44117 • Optional left precision
- 
- 44118 • Optional right precision
- 
- 44119 • A required conversion specifier character that determines the conversion to be performed

44120 **Flags**

44121 One or more of the following optional flags can be specified to control the conversion:

- 44122 =
- f*
- An '=' followed by a single character
- f*
- which is used as the numeric fill character. In |
- 
- 44123 order to work with precision or width counts, the fill character shall be a single byte |
- 
- 44124 character; if not, the behavior is undefined. The default numeric fill character is the |
- 
- 44125 <space>. This flag does not affect field width filling which always uses the <space>. |
- 
- 44126 This flag is ignored unless a left precision (see below) is specified.
- 
- 44127 ^ Do not format the currency amount with grouping characters. The default is to insert
- 
- 44128 the grouping characters if defined for the current locale.
- 
- 44129 + or ( Specify the style of representing positive and negative currency amounts. Only one of
- 
- 44130 '+' or '(' may be specified. If '+' is specified, the locale's equivalent of '+' and '-'
- 
- 44131 are used (for example, in the U.S., the empty string if positive and '-' if negative). If
- 
- 44132 '(' is specified, negative amounts are enclosed within parentheses. If neither flag is
- 
- 44133 specified, the '+' style is used.
- 
- 44134 ! Suppress the currency symbol from the output conversion.
- 
- 44135 - Specify the alignment. If this flag is present the result of the conversion is left-justified |
- 
- 44136 (padded to the right) rather than right-justified. This flag shall be ignored unless a field |
- 
- 44137 width (see below) is specified. |

44138 **Field Width**

44139 *w* A decimal digit string *w* specifying a minimum field width in bytes in which the result  
 44140 of the conversion is right-justified (or left-justified if the flag '-' is specified). The  
 44141 default is 0.

44142 **Left Precision**

44143 *#n* A '#' followed by a decimal digit string *n* specifying a maximum number of digits  
 44144 expected to be formatted to the left of the radix character. This option can be used to  
 44145 keep the formatted output from multiple calls to the *strfmon()* function aligned in the  
 44146 same columns. It can also be used to fill unused positions with a special character as in  
 44147 "\$\*\*\*123.45". This option causes an amount to be formatted as if it has the number  
 44148 of digits specified by *n*. If more than *n* digit positions are required, this conversion  
 44149 specification is ignored. Digit positions in excess of those actually required are filled  
 44150 with the numeric fill character (see the *=f* flag above).

44151 If grouping has not been suppressed with the '^' flag, and it is defined for the current  
 44152 locale, grouping separators are inserted before the fill characters (if any) are added.  
 44153 Grouping separators are not applied to fill characters even if the fill character is a digit.

44154 To ensure alignment, any characters appearing before or after the number in the  
 44155 formatted output such as currency or sign symbols are padded as necessary with  
 44156 <space>s to make their positive and negative formats an equal length.

44157 **Right Precision**

44158 *.p* A period followed by a decimal digit string *p* specifying the number of digits after the  
 44159 radix character. If the value of the right precision *p* is 0, no radix character appears. If a  
 44160 right precision is not included, a default specified by the current locale is used. The  
 44161 amount being formatted is rounded to the specified number of digits prior to  
 44162 formatting.

44163 **Conversion Specifier Characters**

44164 The conversion specifier characters and their meanings are:

44165 *i* The **double** argument is formatted according to the locale's international currency |  
 44166 format (for example, in the U.S.: USD 1,234.56). If the argument is ±Inf or NaN, the |  
 44167 result of the conversion is unspecified. |

44168 *n* The **double** argument is formatted according to the locale's national currency format |  
 44169 (for example, in the U.S.: \$1,234.56). If the argument is ±Inf or NaN, the result of the |  
 44170 conversion is unspecified. |

44171 *%* Convert to a '%'; no argument is converted. The entire conversion specification shall  
 44172 be %%.

44173 **Locale Information**

44174 The *LC\_MONETARY* category of the program's locale affects the behavior of this function  
 44175 including the monetary radix character (which may be different from the numeric radix  
 44176 character affected by the *LC\_NUMERIC* category), the grouping separator, the currency  
 44177 symbols, and formats. The international currency symbol should be conformant with the  
 44178 ISO 4217:1995 standard.

44179 If the value of *maxsize* is greater than {SSIZE\_MAX}, the result is implementation-defined.

44180 **RETURN VALUE**

44181 If the total number of resulting bytes including the terminating null byte is not more than  
44182 *maxsize*, *strfmon()* shall return the number of bytes placed into the array pointed to by *s*, not  
44183 including the terminating null byte. Otherwise, -1 shall be returned, the contents of the array are |  
44184 unspecified, and *errno* shall be set to indicate the error. |

44185 **ERRORS**

44186 The *strfmon()* function shall fail if:

44187 [E2BIG] Conversion stopped due to lack of space in the buffer.

44188 **EXAMPLES**

44189 Given a locale for the U.S. and the values 123.45, -123.45, and 3456.781:

44190

44191

44192

44193

44194

44195

44196

44197

44198

44199

44200

44201

44202

44203

44204

44205

44206

44207

44208

44209

44210

44211

44212

44213

44214

44215

44216

44217

44218

44219

44220

44221

44222

44223

44224

44225

44226

44227

Conversion Specification	Output	Comments
%n	\$123.45 -\$123.45 \$3,456.78	Default formatting.
%11n	\$123.45 -\$123.45 \$3,456.78	Right align within an 11 character field.
%#5n	\$ 123.45 -\$ 123.45 \$ 3,456.78	Aligned columns for values up to 99,999.
%=*#5n	\$***123.45 -\$***123.45 \$*3,456.78	Specify a fill character.
%=0#5n	\$000123.45 -\$000123.45 \$03,456.78	Fill characters do not use grouping even if the fill character is a digit.
%^#5n	\$ 123.45 -\$ 123.45 \$ 3456.78	Disable the grouping separator.
%^#5.0n	\$ 123 -\$ 123 \$ 3457	Round off to whole units.
%^#5.4n	\$ 123.4500 -\$ 123.4500 \$ 3456.7810	Increase the precision.
%(#5n	\$ 123.45 (\$ 123.45) \$ 3,456.78	Use an alternative pos/neg style.
%(!#5n	123.45 ( 123.45) 3,456.78	Disable the currency symbol.
%-14#5.4n	\$ 123.4500 -\$ 123.4500 \$3,456.7810	Left-justify the output.
%14#5.4n	\$ 123.4500 -\$ 123.4500 \$3,456.7810	Corresponding right-justified output.

44228 **APPLICATION USAGE**

44229 None.

44230 **RATIONALE**

44231 None.



44232 **FUTURE DIRECTIONS**

44233           Lowercase conversion characters are reserved for future standards use and uppercase for  
44234           implementation-defined use.

44235 **SEE ALSO**

44236           *localeconv()*, the Base Definitions volume of IEEE Std 1003.1-200x, <**monetary.h**>

44237 **CHANGE HISTORY**

44238           First released in Issue 4.

44239 **Issue 5**

44240           Moved from ENHANCED I18N to BASE.

44241           The [ENOSYS] error is removed.

44242           A sentence is added to the DESCRIPTION warning about values of *maxsize* that are greater than

44243           {SSIZE\_MAX}.

44244 **Issue 6**

44245           The DESCRIPTION is updated to avoid use of the term “must” for application requirements.

44246           The **restrict** keyword is added to the *strfmon()* prototype for alignment with the

44247           ISO/IEC 9899:1999 standard.

## 44248 NAME

44249 strftime — convert date and time to a string

## 44250 SYNOPSIS

44251 #include &lt;time.h&gt;

```
44252     size_t strftime(char *restrict s, size_t maxsize,
44253                   const char *restrict format, const struct tm *restrict timeptr);
```

## 44254 DESCRIPTION

44255 cx The functionality described on this reference page is aligned with the ISO C standard. Any  
 44256 conflict between the requirements described here and the ISO C standard is unintentional. This  
 44257 volume of IEEE Std 1003.1-200x defers to the ISO C standard.

44258 The *strftime()* function shall place bytes into the array pointed to by *s* as controlled by the string  
 44259 pointed to by *format*. The format is a character string, beginning and ending in its initial shift  
 44260 state, if any. The *format* string consists of zero or more conversion specifications and ordinary  
 44261 characters. A conversion specification consists of a '%' character, possibly followed by an E or O  
 44262 modifier, and a terminating conversion specifier character that determines the conversion  
 44263 specification's behavior. All ordinary characters (including the terminating null byte) are copied  
 44264 unchanged into the array. If copying takes place between objects that overlap, the behavior is  
 44265 undefined. No more than *maxsize* bytes are placed into the array. Each conversion specifier is  
 44266 replaced by appropriate characters as described in the following list. The appropriate characters  
 44267 are determined using the *LC\_TIME* category of the current locale and by the values of zero or  
 44268 more members of the broken-down time structure pointed to by *timeptr*, as specified in brackets  
 44269 in the description. If any of the specified values are outside the normal range, the characters  
 44270 stored are unspecified.

44271 cx Local timezone information is used as though *strftime()* called *tzset()*.

44272 The following conversion specifications are supported:

44273	%a	Replaced by the locale's abbreviated weekday name. [ <i>tm_wday</i> ]
44274	%A	Replaced by the locale's full weekday name. [ <i>tm_wday</i> ]
44275	%b	Replaced by the locale's abbreviated month name. [ <i>tm_mon</i> ]
44276	%B	Replaced by the locale's full month name. [ <i>tm_mon</i> ]
44277	%c	Replaced by the locale's appropriate date and time representation. (See the Base 44278 Definitions volume of IEEE Std 1003.1-200x, < <b>time.h</b> >.)
44279	%C	Replaced by the year divided by 100 and truncated to an integer, as a decimal number 44280 [00,99]. [ <i>tm_year</i> ]
44281	%d	Replaced by the day of the month as a decimal number [01,31]. [ <i>tm_mday</i> ]
44282	%D	Equivalent to %m/%d/%y. [ <i>tm_mon</i> , <i>tm_mday</i> , <i>tm_year</i> ]
44283	%e	Replaced by the day of the month as a decimal number [1,31]; a single digit is preceded 44284 by a space. [ <i>tm_mday</i> ]
44285	%F	Equivalent to %Y-%m-%d (the ISO 8601:2000 standard date format). [ <i>tm_year</i> , <i>tm_mon</i> , 44286 <i>tm_mday</i> ]
44287	%g	Replaced by the last 2 digits of the week-based year (see below) as a decimal number 44288 [00,99]. [ <i>tm_year</i> , <i>tm_wday</i> , <i>tm_yday</i> ]
44289	%G	Replaced by the week-based year (see below) as a decimal number (for example, 1977). 44290 [ <i>tm_year</i> , <i>tm_wday</i> , <i>tm_yday</i> ]

44291	%h	Equivalent to %b. [ <i>tm_mon</i> ]	
44292	%H	Replaced by the hour (24-hour clock) as a decimal number [00,23]. [ <i>tm_hour</i> ]	
44293	%I	Replaced by the hour (12-hour clock) as a decimal number [01,12]. [ <i>tm_hour</i> ]	
44294	%j	Replaced by the day of the year as a decimal number [001,366]. [ <i>tm_yday</i> ]	
44295	%m	Replaced by the month as a decimal number [01,12]. [ <i>tm_mon</i> ]	
44296	%M	Replaced by the minute as a decimal number [00,59]. [ <i>tm_min</i> ]	
44297	%n	Replaced by a <newline>.	
44298	%p	Replaced by the locale's equivalent of either a.m. or p.m. [ <i>tm_hour</i> ]	
44299 CX	%r	Replaced by the time in a.m. and p.m. notation; in the POSIX locale this shall be	
44300		equivalent to %I:%M:%S %p. [ <i>tm_hour</i> , <i>tm_min</i> , <i>tm_sec</i> ]	
44301	%R	Replaced by the time in 24 hour notation (%H:%M). [ <i>tm_hour</i> , <i>tm_min</i> ]	
44302	%S	Replaced by the second as a decimal number [00,60]. [ <i>tm_sec</i> ]	
44303	%t	Replaced by a <tab>.	
44304	%T	Replaced by the time (%H:%M:%S). [ <i>tm_hour</i> , <i>tm_min</i> , <i>tm_sec</i> ]	
44305	%u	Replaced by the weekday as a decimal number [1,7], with 1 representing Monday.	
44306		[ <i>tm_wday</i> ]	
44307	%U	Replaced by the week number of the year as a decimal number [00,53]. The first	
44308		Sunday of January is the first day of week 1; days in the new year before this are in	
44309		week 0. [ <i>tm_year</i> , <i>tm_wday</i> , <i>tm_yday</i> ]	
44310	%V	Replaced by the week number of the year (Monday as the first day of the week) as a	
44311		decimal number [01,53]. If the week containing 1 January has four or more days in the	
44312		new year, then it is considered week 1. Otherwise, it is the last week of the previous	
44313		year, and the next week is week 1. Both January 4th and the first Thursday of January	
44314		are always in week 1. [ <i>tm_year</i> , <i>tm_wday</i> , <i>tm_yday</i> ]	
44315	%w	Replaced by the weekday as a decimal number [0,6], with 0 representing Sunday.	
44316		[ <i>tm_wday</i> ]	
44317	%W	Replaced by the week number of the year as a decimal number [00,53]. The first	
44318		Monday of January is the first day of week 1; days in the new year before this are in	
44319		week 0. [ <i>tm_year</i> , <i>tm_wday</i> , <i>tm_yday</i> ]	
44320	%x	Replaced by the locale's appropriate date representation. (See the Base Definitions	
44321		volume of IEEE Std 1003.1-200x, <time.h>.)	
44322	%X	Replaced by the locale's appropriate time representation. (See the Base Definitions	
44323		volume of IEEE Std 1003.1-200x, <time.h>.)	
44324	%y	Replaced by the last two digits of the year as a decimal number [00,99]. [ <i>tm_year</i> ]	
44325	%Y	Replaced by the year as a decimal number (for example, 1997). [ <i>tm_year</i> ]	
44326	%z	Replaced by the offset from UTC in the ISO 8601:2000 standard format (+hhmm or	
44327		-hhmm), or by no characters if no timezone is determinable. For example, "-0430"	
44328 CX		means 4 hours 30 minutes behind UTC (west of Greenwich). If <i>tm_isdst</i> is zero, the	
44329		standard time offset is used. If <i>tm_isdst</i> is greater than zero, the daylight savings time	
44330		offset is used. If <i>tm_isdst</i> is negative, no characters are returned. [ <i>tm_isdst</i> ]	

44331	%Z	Replaced by the timezone name or abbreviation, or by no bytes if no timezone information exists. [ <i>tm_isdst</i> ]	
44332			
44333	%%	Replaced by %.	
44334		If a conversion specification does not correspond to any of the above, the behavior is undefined.	
44335		<b>Modified Conversion Specifiers</b>	
44336		Some conversion specifiers can be modified by the <i>E</i> or <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 conversion specifier. If the alternative format or specification does not exist for the current locale, (see ERA in the Base Definitions volume of IEEE Std 1003.1-200x, Section 7.3.5, LC_TIME) the behavior shall be as if the unmodified conversion specification were used.	
44337			
44338			
44339			
44340			
44341	%Ec	Replaced by the locale's alternative appropriate date and time representation.	
44342	%EC	Replaced by the name of the base year (period) in the locale's alternative representation.	
44343			
44344	%Ex	Replaced by the locale's alternative date representation.	
44345	%EX	Replaced by the locale's alternative time representation.	
44346	%Ey	Replaced by the offset from %EC (year only) in the locale's alternative representation.	
44347	%EY	Replaced by the full alternative year representation.	
44348	%Od	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.	
44349			
44350			
44351	%Oe	Replaced by the day of the month, using the locale's alternative numeric symbols, filled as needed with leading spaces.	
44352			
44353	%OH	Replaced by the hour (24-hour clock) using the locale's alternative numeric symbols.	
44354	%OI	Replaced by the hour (12-hour clock) using the locale's alternative numeric symbols.	
44355	%Om	Replaced by the month using the locale's alternative numeric symbols.	
44356	%OM	Replaced by the minutes using the locale's alternative numeric symbols.	
44357	%OS	Replaced by the seconds using the locale's alternative numeric symbols.	
44358	%Ou	Replaced by the weekday as a number in the locale's alternative representation (Monday=1).	
44359			
44360	%OU	Replaced by the week number of the year (Sunday as the first day of the week, rules corresponding to %U) using the locale's alternative numeric symbols.	
44361			
44362	%OV	Replaced by the week number of the year (Monday as the first day of the week, rules corresponding to %V) using the locale's alternative numeric symbols.	
44363			
44364	%Ow	Replaced by the number of the weekday (Sunday=0) using the locale's alternative numeric symbols.	
44365			
44366	%OW	Replaced by the week number of the year (Monday as the first day of the week) using the locale's alternative numeric symbols.	
44367			
44368	%Oy	Replaced by the year (offset from %C) using the locale's alternative numeric symbols.	
44369	%g, %G, and %V	give values according to the ISO 8601:2000 standard week-based year. In this system, weeks begin on a Monday and week 1 of the year is the week that includes January 4th,	
44370			

44371 which is also the week that includes the first Thursday of the year, and is also the first week that  
 44372 contains at least four days in the year. If the first Monday of January is the 2nd, 3rd, or 4th, the  
 44373 preceding days are part of the last week of the preceding year; thus, for Saturday 2nd January  
 44374 1999, %G is replaced by 1998 and %V is replaced by 53. If December 29th, 30th, or 31st is a  
 44375 Monday, it and any following days are part of week 1 of the following year. Thus, for Tuesday  
 44376 30th December 1997, %G is replaced by 1998 and %V is replaced by 01.

44377 If a conversion specifier is not one of the above, the behavior is undefined. |

#### 44378 RETURN VALUE

44379 If the total number of resulting bytes including the terminating null byte is not more than  
 44380 *maxsize*, *strptime()* shall return the number of bytes placed into the array pointed to by *s*, not  
 44381 including the terminating null byte. Otherwise, 0 shall be returned and the contents of the array  
 44382 are unspecified. |

#### 44383 ERRORS

44384 No errors are defined.

#### 44385 EXAMPLES

##### 44386 Getting a Localized Date String

44387 The following example first sets the locale to the user's default. The locale information will be  
 44388 used in the *nl\_langinfo()* and *strptime()* functions. The *nl\_langinfo()* function returns the localized  
 44389 date string which specifies how the date is laid out. The *strptime()* function takes this information  
 44390 and, using the **tm** structure for values, places the date and time information into *datestring*.

```
44391 #include <time.h>
44392 #include <locale.h>
44393 #include <langinfo.h>
44394 ...
44395 struct tm *tm;
44396 char datestring[256];
44397 ...
44398 setlocale (LC_ALL, "");
44399 ...
44400 strptime (datestring, sizeof(datestring), nl_langinfo (D_T_FMT), tm);
44401 ...
```

#### 44402 APPLICATION USAGE

44403 The range of values for %S is [00,60] rather than [00,59] to allow for the occasional leap second.

44404 Some of the conversion specifications are duplicates of others. They are included for |  
 44405 compatibility with *nl\_cxtime()* and *nl\_ascxtime()*, which were published in Issue 2. |

44406 Applications should use %Y (4-digit years) in preference to %y (2-digit years).

44407 In the C locale, the E and O modifiers are ignored and the replacement strings for the following  
 44408 specifiers are:

44409	%a	The first three characters of %A.
44410	%A	One of Sunday, Monday, ..., Saturday.
44411	%b	The first three characters of %B.
44412	%B	One of January, February, ..., December.
44413	%c	Equivalent to %a %b %e %T %Y.

44414	%p	One of AM or PM.	
44415	%r	Equivalent to %I:%M:%S %p.	
44416	%x	Equivalent to %m/%d/%y.	
44417	%X	Equivalent to %T.	
44418	%Z	Implementation-defined.	
44419	<b>RATIONALE</b>		
44420		None.	
44421	<b>FUTURE DIRECTIONS</b>		
44422		None.	
44423	<b>SEE ALSO</b>		
44424		<i>asctime()</i> , <i>clock()</i> , <i>ctime()</i> , <i>difftime()</i> , <i>getdate()</i> , <i>gmtime()</i> , <i>localtime()</i> , <i>mktime()</i> , <i>strptime()</i> , <i>time()</i> ,	
44425		<i>tzset()</i> , <i>utime()</i> , the Base Definitions volume of IEEE Std 1003.1-200x, < <b>time.h</b> >	
44426	<b>CHANGE HISTORY</b>		
44427		First released in Issue 3.	
44428	<b>Issue 5</b>		
44429		The description of %OV is changed to be consistent with %v and defines Monday as the first day	
44430		of the week.	
44431		The description of %Oy is clarified.	
44432	<b>Issue 6</b>		
44433		Extensions beyond the ISO C standard are now marked.	
44434		The Open Group Corrigendum U033/8 is applied. The %v conversion specifier is changed from	
44435		“Otherwise, it is week 53 of the previous year, and the next week is week 1” to “Otherwise, it is	
44436		the last week of the previous year, and the next week is week 1”.	
44437		The following new requirements on POSIX implementations derive from alignment with the	
44438		Single UNIX Specification:	
44439		<ul style="list-style-type: none"> <li>• The %C, %D, %e, %h, %n, %r, %R, %t, and %T conversion specifiers are added.</li> </ul>	
44440		<ul style="list-style-type: none"> <li>• The modified conversion specifiers are added for consistency with the ISO POSIX-2 standard</li> </ul>	
44441		<i>date</i> utility.	
44442		The following changes are made for alignment with the ISO/IEC 9899: 1999 standard:	
44443		<ul style="list-style-type: none"> <li>• The <i>strptime()</i> prototype is updated.</li> </ul>	
44444		<ul style="list-style-type: none"> <li>• The DESCRIPTION is extensively revised.</li> </ul>	
44445		<ul style="list-style-type: none"> <li>• The %z conversion specifier is added.</li> </ul>	
44446		A new example is added.	

44447 **NAME**

44448            strlen — get string length

44449 **SYNOPSIS**

44450            #include &lt;string.h&gt;

44451            size\_t strlen(const char \*s);

44452 **DESCRIPTION**

44453 cx        The functionality described on this reference page is aligned with the ISO C standard. Any  
 44454            conflict between the requirements described here and the ISO C standard is unintentional. This  
 44455            volume of IEEE Std 1003.1-200x defers to the ISO C standard.

44456            The *strlen()* function shall compute the number of bytes in the string to which *s* points, not  
 44457            including the terminating null byte.

44458 **RETURN VALUE**

44459            The *strlen()* function shall return the length of *s*; no return value shall be reserved to indicate an  
 44460            error.

44461 **ERRORS**

44462            No errors are defined.

44463 **EXAMPLES**44464            **Getting String Lengths**

44465            The following example sets the maximum length of *key* and *data* by using *strlen()* to get the  
 44466            lengths of those strings.

```
44467            #include <string.h>
44468            ...
44469            struct element {
44470                char *key;
44471                char *data;
44472            };
44473            ...
44474            char *key, *data;
44475            int len;

44476            *keylength = *datalength = 0;
44477            ...
44478            if ((len = strlen(key)) > *keylength)
44479                *keylength = len;
44480            if ((len = strlen(data)) > *datalength)
44481                *datalength = len;
44482            ...
```

44483 **APPLICATION USAGE**

44484            None.

44485 **RATIONALE**

44486            None.

44487 **FUTURE DIRECTIONS**

44488            None.

44489 **SEE ALSO**44490           The Base Definitions volume of IEEE Std 1003.1-200x, <**string.h**>44491 **CHANGE HISTORY**

44492           First released in Issue 1. Derived from Issue 1 of the SVID.

44493 **Issue 5**44494           The RETURN VALUE section is updated to indicate that *strlen()* returns the length of *s*, and not  
44495           *s* itself as was previously stated.



44496 **NAME**

44497       strncasecmp — case-insensitive string comparison

44498 **SYNOPSIS**

44499 xSI       #include &lt;strings.h&gt;

44500       int strncasecmp(const char \*s1, const char \*s2, size\_t n);

44501

44502 **DESCRIPTION**44503       Refer to *strcasecmp()*.

44504 **NAME**

44505           strncat — concatenate a string with part of another

44506 **SYNOPSIS**

44507           #include <string.h>

44508           char \*strncat(char \*restrict *s1*, const char \*restrict *s2*, size\_t *n*);

44509 **DESCRIPTION**

44510 **CX**       The functionality described on this reference page is aligned with the ISO C standard. Any  
44511 conflict between the requirements described here and the ISO C standard is unintentional. This  
44512 volume of IEEE Std 1003.1-200x defers to the ISO C standard.

44513       The *strncat()* function shall append not more than *n* bytes (a null byte and bytes that follow it  
44514 are not appended) from the array pointed to by *s2* to the end of the string pointed to by *s1*. The  
44515 initial byte of *s2* overwrites the null byte at the end of *s1*. A terminating null byte is always  
44516 appended to the result. If copying takes place between objects that overlap, the behavior is  
44517 undefined.

44518 **RETURN VALUE**

44519       The *strncat()* function shall return *s1*; no return value shall be reserved to indicate an error.

44520 **ERRORS**

44521       No errors are defined.

44522 **EXAMPLES**

44523       None.

44524 **APPLICATION USAGE**

44525       None.

44526 **RATIONALE**

44527       None.

44528 **FUTURE DIRECTIONS**

44529       None.

44530 **SEE ALSO**

44531       *strcat()*, the Base Definitions volume of IEEE Std 1003.1-200x, <**string.h**>

44532 **CHANGE HISTORY**

44533       First released in Issue 1. Derived from Issue 1 of the SVID.

44534 **Issue 6**

44535       The *strncat()* prototype is updated for alignment with the ISO/IEC 9899:1999 standard.

44536 **NAME**

44537 strncmp — compare part of two strings

44538 **SYNOPSIS**

44539 #include &lt;string.h&gt;

44540 int strncmp(const char \*s1, const char \*s2, size\_t n);

44541 **DESCRIPTION**

44542 cx The functionality described on this reference page is aligned with the ISO C standard. Any  
44543 conflict between the requirements described here and the ISO C standard is unintentional. This  
44544 volume of IEEE Std 1003.1-200x defers to the ISO C standard.

44545 The *strncmp()* function shall compare not more than *n* bytes (bytes that follow a null byte are not  
44546 compared) from the array pointed to by *s1* to the array pointed to by *s2*.

44547 The sign of a non-zero return value is determined by the sign of the difference between the  
44548 values of the first pair of bytes (both interpreted as type **unsigned char**) that differ in the strings  
44549 being compared.

44550 **RETURN VALUE**

44551 Upon successful completion, *strncmp()* shall return an integer greater than, equal to, or less than  
44552 0, if the possibly null-terminated array pointed to by *s1* is greater than, equal to, or less than the  
44553 possibly null-terminated array pointed to by *s2* respectively.

44554 **ERRORS**

44555 No errors are defined.

44556 **EXAMPLES**

44557 None.

44558 **APPLICATION USAGE**

44559 None.

44560 **RATIONALE**

44561 None.

44562 **FUTURE DIRECTIONS**

44563 None.

44564 **SEE ALSO**44565 *strcmp()*, the Base Definitions volume of IEEE Std 1003.1-200x, <**string.h**>44566 **CHANGE HISTORY**

44567 First released in Issue 1. Derived from Issue 1 of the SVID.

44568 **Issue 6**

44569 Extensions beyond the ISO C standard are now marked.

44570 **NAME**

44571       strncpy — copy part of a string

44572 **SYNOPSIS**

44573       #include <string.h>

44574       char \*strncpy(char \*restrict *s1*, const char \*restrict *s2*, size\_t *n*);

44575 **DESCRIPTION**

44576 **CX**       The functionality described on this reference page is aligned with the ISO C standard. Any  
44577 conflict between the requirements described here and the ISO C standard is unintentional. This  
44578 volume of IEEE Std 1003.1-200x defers to the ISO C standard.

44579       The *strncpy()* function shall copy not more than *n* bytes (bytes that follow a null byte are not  
44580 copied) from the array pointed to by *s2* to the array pointed to by *s1*. If copying takes place  
44581 between objects that overlap, the behavior is undefined.

44582       If the array pointed to by *s2* is a string that is shorter than *n* bytes, null bytes shall be appended  
44583 to the copy in the array pointed to by *s1*, until *n* bytes in all are written.

44584 **RETURN VALUE**

44585       The *strncpy()* function shall return *s1*; no return value is reserved to indicate an error.

44586 **ERRORS**

44587       No errors are defined.

44588 **EXAMPLES**

44589       None.

44590 **APPLICATION USAGE**

44591       Character movement is performed differently in different implementations. Thus, overlapping  
44592 moves may yield surprises.

44593       If there is no null byte in the first *n* bytes of the array pointed to by *s2*, the result is not null-  
44594 terminated.

44595 **RATIONALE**

44596       None.

44597 **FUTURE DIRECTIONS**

44598       None.

44599 **SEE ALSO**

44600       *strcpy()*, the Base Definitions volume of IEEE Std 1003.1-200x, <**string.h**>

44601 **CHANGE HISTORY**

44602       First released in Issue 1. Derived from Issue 1 of the SVID.

44603 **Issue 6**

44604       The *strncpy()* prototype is updated for alignment with the ISO/IEC 9899:1999 standard.

44605 **NAME**

44606 strpbrk — scan string for byte

44607 **SYNOPSIS**

44608 #include &lt;string.h&gt;

44609 char \*strpbrk(const char \*s1, const char \*s2);

44610 **DESCRIPTION**

44611 cx The functionality described on this reference page is aligned with the ISO C standard. Any  
44612 conflict between the requirements described here and the ISO C standard is unintentional. This  
44613 volume of IEEE Std 1003.1-200x defers to the ISO C standard.

44614 The *strpbrk()* function shall locate the first occurrence in the string pointed to by *s1* of any byte  
44615 from the string pointed to by *s2*.

44616 **RETURN VALUE**

44617 Upon successful completion, *strpbrk()* shall return a pointer to the byte or a null pointer if no  
44618 byte from *s2* occurs in *s1*.

44619 **ERRORS**

44620 No errors are defined.

44621 **EXAMPLES**

44622 None.

44623 **APPLICATION USAGE**

44624 None.

44625 **RATIONALE**

44626 None.

44627 **FUTURE DIRECTIONS**

44628 None.

44629 **SEE ALSO**44630 *strchr()*, *strchr()*, the Base Definitions volume of IEEE Std 1003.1-200x, <string.h>44631 **CHANGE HISTORY**

44632 First released in Issue 1. Derived from Issue 1 of the SVID.

## 44633 NAME

44634 strptime — date and time conversion

## 44635 SYNOPSIS

44636 XSI #include &lt;time.h&gt;

```
44637 char *strptime(const char *restrict buf, const char *restrict format,
44638               struct tm *restrict tm);
44639
```

## 44640 DESCRIPTION

44641 The *strptime()* function shall convert the character string pointed to by *buf* to values which are  
 44642 stored in the **tm** structure pointed to by *tm*, using the format specified by *format*.

44643 The *format* is composed of zero or more directives. Each directive is composed of one of the  
 44644 following: one or more white-space characters (as specified by *isspace()*); an ordinary character  
 44645 (neither '%' nor a white-space character); or a conversion specification. Each conversion  
 44646 specification is composed of a '%' character followed by a conversion character which specifies  
 44647 the replacement required. The application shall ensure that there is white-space or other non-  
 44648 alphanumeric characters between any two conversion specifications. The following conversion  
 44649 specifications are supported:

44650	%a	The day of the week, using the locale's weekday names; either the abbreviated or full name may be specified.	
44651			
44652	%A	Equivalent to %a.	
44653	%b	The month, using the locale's month names; either the abbreviated or full name may be specified.	
44654			
44655	%B	Equivalent to %b.	
44656	%c	Replaced by the locale's appropriate date and time representation.	
44657	%C	The century number [0,99]; leading zeros are permitted but not required.	
44658	%d	The day of the month [1,31]; leading zeros are permitted but not required.	
44659	%D	The date as %m/%d/%y.	
44660	%e	Equivalent to %d.	
44661	%h	Equivalent to %b.	
44662	%H	The hour (24-hour clock) [0,23]; leading zeros are permitted but not required.	
44663	%I	The hour (12-hour clock) [1,12]; leading zeros are permitted but not required.	
44664	%j	The day number of the year [1,366]; leading zeros are permitted but not required.	
44665	%m	The month number [1,12]; leading zeros are permitted but not required.	
44666	%M	The minute [0,59]; leading zeros are permitted but not required.	
44667	%n	Any white space.	
44668	%p	The locale's equivalent of a.m or p.m.	
44669	%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 shall be equivalent	
44670		to %I:%M:%S %p.	
44671			
44672	%R	The time as %H:%M.	

44673	%S	The seconds [0,60]; leading zeros are permitted but not required.
44674	%t	Any white space.
44675	%T	The time as %H:%M:%S.
44676	%U	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.
44678	%w	The weekday as a decimal number [0,6], with 0 representing Sunday; leading zeros are permitted but not required.
44680	%W	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.
44682	%x	The date, using the locale's date format.
44683	%X	The time, using the locale's time format.
44684	%y	The year within century. When a century is not otherwise specified, values in the range [69,99] shall refer to years 1969 to 1999 inclusive, and values in the range [00,68] shall refer to years 2000 to 2068 inclusive; leading zeros shall be permitted but shall not be required.
44688	<b>Note:</b>	It is expected that in a future version of IEEE Std 1003.1-200x the default century inferred from a 2-digit year will change. (This would apply to all commands accepting a 2-digit year as input.)
44691	%Y	The year, including the century (for example, 1988).
44692	%%	Replaced by %.

#### 44693 **Modified Conversion Specifiers**

44694		Some conversion specifiers can be modified by the <b>E</b> and <b>O</b> modifier characters to indicate that an alternative format or specification should be used rather than the one normally used by the unmodified conversion specifier. If the alternative format or specification does not exist in the current locale, the behavior shall be as if the unmodified conversion specification were used.
44698	%Ec	The locale's alternative appropriate date and time representation.
44699	%EC	The name of the base year (period) in the locale's alternative representation.
44700	%Ex	The locale's alternative date representation.
44701	%EX	The locale's alternative time representation.
44702	%Ey	The offset from %EC (year only) in the locale's alternative representation.
44703	%EY	The full alternative year representation.
44704	%Od	The day of the month using the locale's alternative numeric symbols; leading zeros are permitted but not required.
44706	%Oe	Equivalent to %Od.
44707	%OH	The hour (24-hour clock) using the locale's alternative numeric symbols.
44708	%OI	The hour (12-hour clock) using the locale's alternative numeric symbols.
44709	%Om	The month using the locale's alternative numeric symbols.
44710	%OM	The minutes using the locale's alternative numeric symbols.

44711	%OS	The seconds using the locale's alternative numeric symbols.
44712	%OU	The week number of the year (Sunday as the first day of the week) using the locale's alternative numeric symbols.
44713		
44714	%Ow	The number of the weekday (Sunday=0) using the locale's alternative numeric symbols.
44715	%OW	The week number of the year (Monday as the first day of the week) using the locale's alternative numeric symbols.
44716		
44717	%Oy	The year (offset from %C) using the locale's alternative numeric symbols.
44718		A conversion specification composed of white-space characters is executed by scanning input up to the first character that is not white-space (which remains unscanned), or until no more characters can be scanned.
44719		
44720		
44721		A conversion specification that is an ordinary character is executed by scanning the next character from the buffer. If the character scanned from the buffer differs from the one comprising the directive, the directive fails, and the differing and subsequent characters remain unscanned.
44722		
44723		
44724		
44725		A series of conversion specifications composed of %n, %t, white-space characters, or any combination is executed by scanning up to the first character that is not white space (which remains unscanned), or until no more characters can be scanned.
44726		
44727		
44728		Any other conversion specification is executed by scanning characters until a character matching the next directive is scanned, or until no more characters can be scanned. These characters, except the one matching the next directive, are then compared to the locale values associated with the conversion specifier. If a match is found, values for the appropriate <b>tm</b> structure members are set to values corresponding to the locale information. Case is ignored when matching items in <i>buf</i> such as month or weekday names. If no match is found, <i>strptime()</i> fails and no more characters are scanned.
44729		
44730		
44731		
44732		
44733		
44734		
44735	<b>RETURN VALUE</b>	
44736		Upon successful completion, <i>strptime()</i> shall return a pointer to the character following the last character parsed. Otherwise, a null pointer shall be returned.
44737		
44738	<b>ERRORS</b>	
44739		No errors are defined.
44740	<b>EXAMPLES</b>	
44741		None.
44742	<b>APPLICATION USAGE</b>	
44743		Several "equivalent to" formats and the special processing of white-space characters are provided in order to ease the use of identical <i>format</i> strings for <i>strptime()</i> and <i>strptime()</i> .
44744		
44745		Applications should use %Y (4-digit years) in preference to %y (2-digit years).
44746		It is unspecified whether multiple calls to <i>strptime()</i> using the same <b>tm</b> structure will update the current contents of the structure or overwrite all contents of the structure. Conforming applications should make a single call to <i>strptime()</i> with a format and all data needed to completely specify the date and time being converted.
44747		
44748		
44749		
44750	<b>RATIONALE</b>	
44751		None.



44752 **FUTURE DIRECTIONS**

44753           The *strptime()* function is expected to be mandatory in the next version of this volume of  
44754           IEEE Std 1003.1-200x.

44755 **SEE ALSO**

44756           *scanf()*, *strptime()*, *time()*, the Base Definitions volume of IEEE Std 1003.1-200x, <**time.h**>

44757 **CHANGE HISTORY**

44758           First released in Issue 4.

44759 **Issue 5**

44760           Moved from ENHANCED I18N to BASE.

44761           The [ENOSYS] error is removed.

44762           The exact meaning of the %y and %Oy specifiers are clarified in the DESCRIPTION.

44763 **Issue 6**

44764           The Open Group Corrigendum U033/5 is applied. The %r specifier description is reworded.

44765           The DESCRIPTION is updated to avoid use of the term “must” for application requirements.

44766           The **restrict** keyword is added to the *strptime()* prototype for alignment with the  
44767           ISO/IEC 9899:1999 standard.

44768           The Open Group Corrigendum U047/2 is applied.

44769           The DESCRIPTION is updated to use the terms “conversion specifier” and “conversion  
44770           specification” for consistency with *strptime()*.

44771 **NAME**

44772       strchr — string scanning operation

44773 **SYNOPSIS**

44774       #include &lt;string.h&gt;

44775       char \*strchr(const char \*s, int c);

44776 **DESCRIPTION**

44777 CX       The functionality described on this reference page is aligned with the ISO C standard. Any  
44778       conflict between the requirements described here and the ISO C standard is unintentional. This  
44779       volume of IEEE Std 1003.1-200x defers to the ISO C standard.

44780 CX       The *strchr()* function shall locate the last occurrence of *c* (converted to an **unsigned char**) in the  
44781       string pointed to by *s*. The terminating null byte is considered to be part of the string.

44782 **RETURN VALUE**

44783       Upon successful completion, *strchr()* shall return a pointer to the byte or a null pointer if *c* does  
44784       not occur in the string.

44785 **ERRORS**

44786       No errors are defined.

44787 **EXAMPLES**44788       **Finding the Base Name of a File**

44789       The following example uses *strchr()* to get a pointer to the base name of a file. The *strchr()*  
44790       function searches backwards through the name of the file to find the last '/' character in *name*.  
44791       This pointer (plus one) will point to the base name of the file.

```
44792       #include <string.h>
44793       ...
44794       const char *name;
44795       char *basename;
44796       ...
44797       basename = strchr(name, '/') + 1;
44798       ...
```

44799 **APPLICATION USAGE**

44800       None.

44801 **RATIONALE**

44802       None.

44803 **FUTURE DIRECTIONS**

44804       None.

44805 **SEE ALSO**44806       *strchr()*, the Base Definitions volume of IEEE Std 1003.1-200x, <string.h>44807 **CHANGE HISTORY**

44808       First released in Issue 1. Derived from Issue 1 of the SVID.

44809 **NAME**

44810 strspn — get length of a substring

44811 **SYNOPSIS**

44812 #include &lt;string.h&gt;

44813 size\_t strspn(const char \*s1, const char \*s2);

44814 **DESCRIPTION**

44815 cx The functionality described on this reference page is aligned with the ISO C standard. Any  
44816 conflict between the requirements described here and the ISO C standard is unintentional. This  
44817 volume of IEEE Std 1003.1-200x defers to the ISO C standard.

44818 The *strspn()* function shall compute the length (in bytes) of the maximum initial segment of the  
44819 string pointed to by *s1* which consists entirely of bytes from the string pointed to by *s2*.

44820 **RETURN VALUE**

44821 The *strspn()* function shall return the length of *s1*; no return value is reserved to indicate an  
44822 error.

44823 **ERRORS**

44824 No errors are defined.

44825 **EXAMPLES**

44826 None.

44827 **APPLICATION USAGE**

44828 None.

44829 **RATIONALE**

44830 None.

44831 **FUTURE DIRECTIONS**

44832 None.

44833 **SEE ALSO**44834 *strcspn()*, the Base Definitions volume of IEEE Std 1003.1-200x, <string.h>44835 **CHANGE HISTORY**

44836 First released in Issue 1. Derived from Issue 1 of the SVID.

44837 **Issue 5**

44838 The RETURN VALUE section is updated to indicate that *strspn()* returns the length of *s*, and not  
44839 *s* itself as was previously stated.

44840 **NAME**44841        `strstr` — find a substring44842 **SYNOPSIS**44843        `#include <string.h>`44844        `char *strstr(const char *s1, const char *s2);`44845 **DESCRIPTION**

44846 `CX`        The functionality described on this reference page is aligned with the ISO C standard. Any  
44847        conflict between the requirements described here and the ISO C standard is unintentional. This  
44848        volume of IEEE Std 1003.1-200x defers to the ISO C standard.

44849        The `strstr()` function shall locate the first occurrence in the string pointed to by `s1` of the  
44850        sequence of bytes (excluding the terminating null byte) in the string pointed to by `s2`.

44851 **RETURN VALUE**

44852        Upon successful completion, `strstr()` shall return a pointer to the located string or a null pointer  
44853        if the string is not found.

44854        If `s2` points to a string with zero length, the function shall return `s1`.

44855 **ERRORS**

44856        No errors are defined.

44857 **EXAMPLES**

44858        None.

44859 **APPLICATION USAGE**

44860        None.

44861 **RATIONALE**

44862        None.

44863 **FUTURE DIRECTIONS**

44864        None.

44865 **SEE ALSO**

44866        `strchr()`, the Base Definitions volume of IEEE Std 1003.1-200x, `<string.h>`

44867 **CHANGE HISTORY**

44868        First released in Issue 3.

44869        Entry included for alignment with the ANSI C standard.

## 44870 NAME

44871 strtod, strtodf, strtold — convert string to a double-precision number

## 44872 SYNOPSIS

44873 #include &lt;stdlib.h&gt;

44874 double strtod(const char \*restrict *nptr*, char \*\*restrict *endp*);44875 float strtodf(const char \*restrict *nptr*, char \*\*restrict *endp*);44876 long double strtold(const char \*restrict *nptr*, char \*\*restrict *endp*);

## 44877 DESCRIPTION

44878 cx The functionality described on this reference page is aligned with the ISO C standard. Any  
 44879 conflict between the requirements described here and the ISO C standard is unintentional. This  
 44880 volume of IEEE Std 1003.1-200x defers to the ISO C standard.

44881 These functions shall convert the initial portion of the string pointed to by *nptr* to **double**, **float**,  
 44882 and **long double** representation, respectively. First, they decompose the input string into three  
 44883 parts:

- 44884 1. An initial, possibly empty, sequence of white-space characters (as specified by *isspace()*)
- 44885 2. A subject sequence interpreted as a floating-point constant or representing infinity or NaN
- 44886 3. A final string of one or more unrecognized characters, including the terminating null byte  
 44887 of the input string

44888 Then they shall attempt to convert the subject sequence to a floating-point number, and return  
 44889 the result.

44890 The expected form of the subject sequence is an optional plus or minus sign, then one of the  
 44891 following:

- 44892 • A non-empty sequence of decimal digits optionally containing a radix character, then an  
 44893 optional exponent part
- 44894 • A 0x or 0X, then a non-empty sequence of hexadecimal digits optionally containing a radix  
 44895 character, then an optional binary exponent part
- 44896 • One of INF or INFINITY, ignoring case
- 44897 • One of NAN or NAN(*n-char-sequence<sub>opt</sub>*), ignoring case in the NAN part, where:

```
44898 n-char-sequence:
44899     digit
44900     nondigit
44901     n-char-sequence digit
44902     n-char-sequence nondigit
```

44903 The subject sequence is defined as the longest initial subsequence of the input string, starting  
 44904 with the first non-white-space character, that is of the expected form. The subject sequence  
 44905 contains no characters if the input string is not of the expected form.

44906 If the subject sequence has the expected form for a floating-point number, the sequence of  
 44907 characters starting with the first digit or the decimal-point character (whichever occurs first)  
 44908 shall be interpreted as a floating constant of the C language, except that the radix character shall  
 44909 be used in place of a period, and that if neither an exponent part nor a radix character appears in  
 44910 a decimal floating-point number, or if a binary exponent part does not appear in a hexadecimal  
 44911 floating-point number, an exponent part of the appropriate type with value zero is assumed to  
 44912 follow the last digit in the string. If the subject sequence begins with a minus sign, the sequence  
 44913 shall be interpreted as negated. A character sequence INF or INFINITY shall be interpreted as an

44914 infinity, if representable in the return type, else as if it were a floating constant that is too large |  
 44915 for the range of the return type. A character sequence NAN or NAN(*n-char-sequence<sub>opt</sub>*) shall be |  
 44916 interpreted as a quiet NaN, if supported in the return type, else as if it were a subject sequence |  
 44917 part that does not have the expected form; the meaning of the *n-char* sequences is |  
 44918 implementation-defined. A pointer to the final string is stored in the object pointed to by *endptr*, |  
 44919 provided that *endptr* is not a null pointer.

44920 If the subject sequence has the hexadecimal form and FLT\_RADIX is a power of 2, the value |  
 44921 resulting from the conversion is correctly rounded.

44922 CX The radix character is defined in the program's locale (category *LC\_NUMERIC*). In the POSIX |  
 44923 locale, or in a locale where the radix character is not defined, the radix character shall default to a |  
 44924 period ('.').

44925 CX In other than the C or POSIX locales, other implementation-defined subject sequences may be |  
 44926 accepted.

44927 If the subject sequence is empty or does not have the expected form, no conversion shall be |  
 44928 performed; the value of *str* is stored in the object pointed to by *endptr*, provided that *endptr* is not |  
 44929 a null pointer.

44930 CX The *strtod()* function shall not change the setting of *errno* if successful.

44931 Since 0 is returned on error and is also a valid return on success, an application wishing to check |  
 44932 for error situations should set *errno* to 0, then call *strtod()*, *strtof()*, or *strtold()*, then check *errno*.

#### 44933 RETURN VALUE

44934 Upon successful completion, these functions shall return the converted value. If no conversion |  
 44935 could be performed, 0 shall be returned, and *errno* may be set to [EINVAL].

44936 If the correct value is outside the range of representable values, HUGE\_VAL, HUGE\_VALF, or |  
 44937 HUGE\_VALL shall be returned (according to the sign of the value), and *errno* shall be set to |  
 44938 [ERANGE].

44939 If the correct value would cause an underflow, a value whose magnitude is no greater than the |  
 44940 smallest normalized positive number in the return type shall be returned and *errno* set to |  
 44941 [ERANGE].

#### 44942 ERRORS

44943 These functions shall fail if:

44944 CX [ERANGE] The value to be returned would cause overflow or underflow.

44945 These functions may fail if:

44946 CX [EINVAL] No conversion could be performed.

#### 44947 EXAMPLES

44948 None.

#### 44949 APPLICATION USAGE

44950 If the subject sequence has the hexadecimal form and FLT\_RADIX is not a power of 2, and the |  
 44951 result is not exactly representable, the result should be one of the two numbers in the |  
 44952 appropriate internal format that are adjacent to the hexadecimal floating source value, with the |  
 44953 extra stipulation that the error should have a correct sign for the current rounding direction. |

44954 If the subject sequence has the decimal form and at most DECIMAL\_DIG (defined in <float.h>) |  
 44955 significant digits, the result should be correctly rounded. If the subject sequence *D* has the |  
 44956 decimal form and more than DECIMAL\_DIG significant digits, consider the two bounding, |  
 44957 adjacent decimal strings *L* and *U*, both having DECIMAL\_DIG significant digits, such that the

44958 values of  $L$ ,  $D$ , and  $U$  satisfy  $L \leq D \leq U$ . The result should be one of the (equal or adjacent)  
 44959 values that would be obtained by correctly rounding  $L$  and  $U$  according to the current rounding  
 44960 direction, with the extra stipulation that the error with respect to  $D$  should have a correct sign  
 44961 for the current rounding direction.

44962 The changes to `strtod()` introduced by the ISO/IEC 9899:1999 standard can alter the behavior of  
 44963 well-formed applications complying with the ISO/IEC 9899:1990 standard and thus earlier  
 44964 versions of IEEE Std 1003.1-200x. One such example would be:

```

44965 int
44966 what_kind_of_number (char *s)
44967 {
44968     char *endp;
44969     double d;
44970     long l;

44971     d = strtod(s, &endp);
44972     if (s != endp && *endp == '\0')
44973         printf("It's a float with value %g\n", d);
44974     else
44975     {
44976         l = strtol(s, &endp, 0);
44977         if (s != endp && *endp == '\0')
44978             printf("It's an integer with value %ld\n", l);
44979         else
44980             return 1;
44981     }
44982     return 0;
44983 }
```

44984 If the function is called with:

```
44985 what_kind_of_number ("0x10")
```

44986 an ISO/IEC 9899:1990 standard-compliant library will result in the function printing:

```
44987 It's an integer with value 16
```

44988 With the ISO/IEC 9899:1999 standard, the result is:

```
44989 It's a float with value 16
```

44990 The change in behavior is due to the inclusion of floating-point numbers in hexadecimal  
 44991 notation without requiring that either a decimal point or the binary exponent be present.

#### 44992 RATIONALE

44993 None.

#### 44994 FUTURE DIRECTIONS

44995 None.

#### 44996 SEE ALSO

44997 `isspace()`, `localeconv()`, `scanf()`, `setlocale()`, `strtol()`, the Base Definitions volume of  
 44998 IEEE Std 1003.1-200x, `<float.h>`, `<stdlib.h>`, the Base Definitions volume of  
 44999 IEEE Std 1003.1-200x, Chapter 7, Locale

45000 **CHANGE HISTORY**

45001 First released in Issue 1. Derived from Issue 1 of the SVID.

45002 **Issue 5**

45003 The DESCRIPTION is updated to indicate that *errno* is not changed if the function is successful.

45004 **Issue 6**

45005 Extensions beyond the ISO C standard are now marked.

45006 The following new requirements on POSIX implementations derive from alignment with the  
45007 Single UNIX Specification:

45008 • In the RETURN VALUE and ERRORS sections, the [EINVAL] optional error condition is  
45009 added if no conversion could be performed.

45010 The following changes are made for alignment with the ISO/IEC 9899:1999 standard:

45011 • The *strtod()* function is updated.

45012 • The *strtof()* and *strtold()* functions are added.

45013 • The DESCRIPTION is extensively revised.

45014 ISO/IEC 9899:1999 standard, Technical Corrigendum No. 1 is incorporated. |



45015 **NAME**

45016        strtod — convert string to a double-precision number

45017 **SYNOPSIS**

45018        #include &lt;stdlib.h&gt;

45019        float strtod(const char \*restrict *nptr*, char \*\*restrict *endptr*);45020 **DESCRIPTION**45021        Refer to *strtod*().

45022 **NAME**

45023 strtoimax, strtoumax — convert string to integer type

45024 **SYNOPSIS**

45025 #include <inttypes.h>

45026 intmax\_t strtoimax(const char \*restrict nptr, char \*\*restrict endptr,  
45027 int base);

45028 uintmax\_t strtoumax(const char \*restrict nptr, char \*\*restrict endptr,  
45029 int base);

45030 **DESCRIPTION**

45031 **CX** The functionality described on this reference page is aligned with the ISO C standard. Any  
45032 conflict between the requirements described here and the ISO C standard is unintentional. This  
45033 volume of IEEE Std 1003.1-200x defers to the ISO C standard.

45034 These functions shall be equivalent to the *strtol()*, *strtoll()*, *strtoul()*, and *strtoull()* functions,  
45035 except that the initial portion of the string shall be converted to **intmax\_t** and **uintmax\_t**  
45036 representation, respectively.

45037 **RETURN VALUE**

45038 These functions shall return the converted value, if any.

45039 If no conversion could be performed, zero shall be returned.

45040 If the correct value is outside the range of representable values, {INTMAX\_MAX},  
45041 {INTMAX\_MIN}, or {UINTMAX\_MAX} shall be returned (according to the return type and sign  
45042 of the value, if any), and *errno* shall be set to [ERANGE].

45043 **ERRORS**

45044 These functions shall fail if:

45045 [ERANGE] The value to be returned is not representable.

45046 These functions may fail if:

45047 [EINVAL] The value of *base* is not supported.

45048 **EXAMPLES**

45049 None.

45050 **APPLICATION USAGE**

45051 None.

45052 **RATIONALE**

45053 None.

45054 **FUTURE DIRECTIONS**

45055 None.

45056 **SEE ALSO**

45057 *strtol()*, *strtoul()*, the Base Definitions volume of IEEE Std 1003.1-200x, <inttypes.h>

45058 **CHANGE HISTORY**

45059 First released in Issue 6. Derived from the ISO/IEC 9899:1999 standard.

## 45060 NAME

45061 strtok, strtok\_r — split string into tokens

## 45062 SYNOPSIS

45063 #include &lt;string.h&gt;

45064 char \*strtok(char \*restrict s1, const char \*restrict s2);

45065 TSF char \*strtok\_r(char \*restrict s, const char \*restrict sep,

45066 char \*\*restrict lasts);

45067

## 45068 DESCRIPTION

45069 CX For *strtok()*: The functionality described on this reference page is aligned with the ISO C  
 45070 standard. Any conflict between the requirements described here and the ISO C standard is  
 45071 unintentional. This volume of IEEE Std 1003.1-200x defers to the ISO C standard.

45072 A sequence of calls to *strtok()* breaks the string pointed to by *s1* into a sequence of tokens, each  
 45073 of which is delimited by a byte from the string pointed to by *s2*. The first call in the sequence has  
 45074 *s1* as its first argument, and is followed by calls with a null pointer as their first argument. The  
 45075 separator string pointed to by *s2* may be different from call to call.

45076 The first call in the sequence searches the string pointed to by *s1* for the first byte that is *not*  
 45077 contained in the current separator string pointed to by *s2*. If no such byte is found, then there  
 45078 are no tokens in the string pointed to by *s1* and *strtok()* shall return a null pointer. If such a byte  
 45079 is found, it is the start of the first token.

45080 The *strtok()* function then searches from there for a byte that *is* contained in the current  
 45081 separator string. If no such byte is found, the current token extends to the end of the string  
 45082 pointed to by *s1*, and subsequent searches for a token shall return a null pointer. If such a byte is  
 45083 found, it is overwritten by a null byte, which terminates the current token. The *strtok()* function  
 45084 saves a pointer to the following byte, from which the next search for a token shall start.

45085 Each subsequent call, with a null pointer as the value of the first argument, starts searching from  
 45086 the saved pointer and behaves as described above.

45087 The implementation shall behave as if no function defined in this volume of  
 45088 IEEE Std 1003.1-200x calls *strtok()*.

45089 CX The *strtok()* function need not be reentrant. A function that is not required to be reentrant is not  
 45090 required to be thread-safe.

45091 TSF The *strtok\_r()* function considers the null-terminated string *s* as a sequence of zero or more text  
 45092 tokens separated by spans of one or more characters from the separator string *sep*. The  
 45093 argument *lasts* points to a user-provided pointer which points to stored information necessary  
 45094 for *strtok\_r()* to continue scanning the same string.

45095 In the first call to *strtok\_r()*, *s* points to a null-terminated string, *sep* to a null-terminated string of  
 45096 separator characters, and the value pointed to by *lasts* is ignored. The *strtok\_r()* function shall  
 45097 return a pointer to the first character of the first token, write a null character into *s* immediately  
 45098 following the returned token, and update the pointer to which *lasts* points.

45099 In subsequent calls, *s* is a NULL pointer and *lasts* shall be unchanged from the previous call so  
 45100 that subsequent calls shall move through the string *s*, returning successive tokens until no  
 45101 tokens remain. The separator string *sep* may be different from call to call. When no token  
 45102 remains in *s*, a NULL pointer shall be returned.

45103 **RETURN VALUE**

45104 Upon successful completion, *strtok()* shall return a pointer to the first byte of a token. Otherwise,  
45105 if there is no token, *strtok()* shall return a null pointer.

45106 TSF The *strtok\_r()* function shall return a pointer to the token found, or a NULL pointer when no  
45107 token is found.

45108 **ERRORS**

45109 No errors are defined.

45110 **EXAMPLES**45111 **Searching for Word Separators**

45112 The following example searches for tokens separated by space characters.

```
45113 #include <string.h>
45114 ...
45115 char *token;
45116 char *line = "LINE TO BE SEPARATED";
45117 char *search = " ";

45118 /* Token will point to "LINE". */
45119 token = strtok(line, search);

45120 /* Token will point to "TO". */
45121 token = strtok(NULL, search);
```

45122 **Breaking a Line**

45123 The following example uses *strtok()* to break a line into two character strings separated by any  
45124 combination of <space>s, <tab>s, or <newline>s.

```
45125 #include <string.h>
45126 ...
45127 struct element {
45128     char *key;
45129     char *data;
45130 };
45131 ...
45132 char line[LINE_MAX];
45133 char *key, *data;
45134 ...
45135 key = strtok(line, " \n");
45136 data = strtok(NULL, " \n");
45137 ...
```

45138 **APPLICATION USAGE**

45139 The *strtok\_r()* function is thread-safe and stores its state in a user-supplied buffer instead of  
45140 possibly using a static data area that may be overwritten by an unrelated call from another  
45141 thread.

45142 **RATIONALE**

45143 The *strtok()* function searches for a separator string within a larger string. It returns a pointer to  
45144 the last substring between separator strings. This function uses static storage to keep track of  
45145 the current string position between calls. The new function, *strtok\_r()*, takes an additional  
45146 argument, *lasts*, to keep track of the current position in the string.

45147 **FUTURE DIRECTIONS**

45148 None.

45149 **SEE ALSO**45150 The Base Definitions volume of IEEE Std 1003.1-200x, <**string.h**>45151 **CHANGE HISTORY**

45152 First released in Issue 1. Derived from Issue 1 of the SVID.

45153 **Issue 5**45154 The *strtok\_r()* function is included for alignment with the POSIX Threads Extension.45155 A note indicating that the *strtok()* function need not be reentrant is added to the DESCRIPTION.45156 **Issue 6**

45157 Extensions beyond the ISO C standard are now marked.

45158 The *strtok\_r()* function is marked as part of the Thread-Safe Functions option.

45159 In the DESCRIPTION, the note about reentrancy is expanded to cover thread-safety.

45160 The APPLICATION USAGE section is updated to include a note on the thread-safe function and

45161 its avoidance of possibly using a static data area.

45162 The **restrict** keyword is added to the *strtok()* and *strtok\_r()* prototypes for alignment with the

45163 ISO/IEC 9899:1999 standard.

45164 **NAME**

45165        strtol, strtoll — convert string to a long integer

45166 **SYNOPSIS**

45167        #include <stdlib.h>

45168        long strtol(const char \*restrict *str*, char \*\*restrict *endptr*, int *base*);

45169        long long strtoll(const char \*restrict *str*, char \*\*restrict *endptr*,

45170            int *base*)

45171 **DESCRIPTION**

45172 **CX**        The functionality described on this reference page is aligned with the ISO C standard. Any  
45173 conflict between the requirements described here and the ISO C standard is unintentional. This  
45174 volume of IEEE Std 1003.1-200x defers to the ISO C standard.

45175        These functions shall convert the initial portion of the string pointed to by *str* to a type **long** and  
45176 **long long** representation, respectively. First, they decompose the input string into three parts:

- 45177            1. An initial, possibly empty, sequence of white-space characters (as specified by *isspace()*)
- 45178            2. A subject sequence interpreted as an integer represented in some radix determined by the  
45179 value of *base*
- 45180            3. A final string of one or more unrecognized characters, including the terminating null byte  
45181 of the input string.

45182        Then they shall attempt to convert the subject sequence to an integer, and return the result. |

45183        If the value of *base* is 0, the expected form of the subject sequence is that of a decimal constant,  
45184 octal constant, or hexadecimal constant, any of which may be preceded by a '+' or '-' sign. A  
45185 decimal constant begins with a non-zero digit, and consists of a sequence of decimal digits. An  
45186 octal constant consists of the prefix '0' optionally followed by a sequence of the digits '0' to  
45187 '7' only. A hexadecimal constant consists of the prefix 0x or 0X followed by a sequence of the  
45188 decimal digits and letters 'a' (or 'A') to 'f' (or 'F') with values 10 to 15 respectively.

45189        If the value of *base* is between 2 and 36, the expected form of the subject sequence is a sequence  
45190 of letters and digits representing an integer with the radix specified by *base*, optionally preceded  
45191 by a '+' or '-' sign. The letters from 'a' (or 'A') to 'z' (or 'Z') inclusive are ascribed the  
45192 values 10 to 35; only letters whose ascribed values are less than that of *base* are permitted. If the  
45193 value of *base* is 16, the characters 0x or 0X may optionally precede the sequence of letters and  
45194 digits, following the sign if present.

45195        The subject sequence is defined as the longest initial subsequence of the input string, starting  
45196 with the first non-white-space character that is of the expected form. The subject sequence shall |  
45197 contain no characters if the input string is empty or consists entirely of white-space characters, |  
45198 or if the first non-white-space character is other than a sign or a permissible letter or digit. |

45199        If the subject sequence has the expected form and the value of *base* is 0, the sequence of |  
45200 characters starting with the first digit shall be interpreted as an integer constant. If the subject |  
45201 sequence has the expected form and the value of *base* is between 2 and 36, it shall be used as the |  
45202 base for conversion, ascribing to each letter its value as given above. If the subject sequence |  
45203 begins with a minus sign, the value resulting from the conversion shall be negated. A pointer to |  
45204 the final string shall be stored in the object pointed to by *endptr*, provided that *endptr* is not a null |  
45205 pointer.

45206 **CX**        In other than the C or POSIX locales, other implementation-defined subject sequences may be  
45207 accepted.

45208 If the subject sequence is empty or does not have the expected form, no conversion is performed;  
 45209 the value of *str* is stored in the object pointed to by *endptr*, provided that *endptr* is not a null  
 45210 pointer.

45211 CX The *strtol()* function shall not change the setting of *errno* if successful.

45212 Since 0, {LONG\_MIN} or {LLONG\_MIN}, and {LONG\_MAX} or {LLONG\_MAX} are returned on  
 45213 error and are also valid returns on success, an application wishing to check for error situations  
 45214 should set *errno* to 0, then call *strtol()* or *strtoll()*, then check *errno*.

#### 45215 RETURN VALUE

45216 Upon successful completion, these functions shall return the converted value, if any. If no  
 45217 CX conversion could be performed, 0 shall be returned and *errno* may be set to [EINVAL].

45218 If the correct value is outside the range of representable values, {LONG\_MIN}, {LONG\_MAX},  
 45219 {LLONG\_MIN} or {LLONG\_MAX} shall be returned (according to the sign of the value), and  
 45220 *errno* set to [ERANGE].

#### 45221 ERRORS

45222 These functions shall fail if:

45223 [ERANGE] The value to be returned is not representable.

45224 These functions may fail if:

45225 CX [EINVAL] The value of *base* is not supported.

#### 45226 EXAMPLES

45227 None.

#### 45228 APPLICATION USAGE

45229 None.

#### 45230 RATIONALE

45231 None.

#### 45232 FUTURE DIRECTIONS

45233 None.

#### 45234 SEE ALSO

45235 *isalpha()*, *scanf()*, *strtod()*, the Base Definitions volume of IEEE Std 1003.1-200x, <stdlib.h>

#### 45236 CHANGE HISTORY

45237 First released in Issue 1. Derived from Issue 1 of the SVID.

#### 45238 Issue 5

45239 The DESCRIPTION is updated to indicate that *errno* is not changed if the function is successful.

#### 45240 Issue 6

45241 Extensions beyond the ISO C standard are now marked.

45242 The following new requirements on POSIX implementations derive from alignment with the  
 45243 Single UNIX Specification:

45244 • In the RETURN VALUE and ERRORS sections, the [EINVAL] optional error condition is  
 45245 added if no conversion could be performed.

45246 The following changes are made for alignment with the ISO/IEC 9899:1999 standard:

45247 • The *strtol()* prototype is updated.

45248 • The *strtoll()* function is added.

45249 **NAME**

45250            **strtold** — convert string to a double-precision number

45251 **SYNOPSIS**

45252            #include <stdlib.h>

45253            long double strtold(const char \*restrict *nptr*, char \*\*restrict *endptr*);

45254 **DESCRIPTION**

45255            Refer to *strtod*().



45256 **NAME**45257            **strtoll** — convert string to a long integer45258 **SYNOPSIS**

45259            #include &lt;stdlib.h&gt;

45260            long long strtoll(const char \*restrict *str*, char \*\*restrict *endptr*,  
45261                            int *base*);45262 **DESCRIPTION**45263            Refer to *strtol*().

## 45264 NAME

45265 strtol, strtoll — convert string to an unsigned long

## 45266 SYNOPSIS

45267 #include &lt;stdlib.h&gt;

45268 unsigned long strtol(const char \*restrict str,  
45269 char \*\*restrict endptr, int base);45270 unsigned long long strtoll(const char \*restrict str,  
45271 char \*\*restrict endptr, int base);

## 45272 DESCRIPTION

45273 cx The functionality described on this reference page is aligned with the ISO C standard. Any  
45274 conflict between the requirements described here and the ISO C standard is unintentional. This  
45275 volume of IEEE Std 1003.1-200x defers to the ISO C standard.45276 These functions shall convert the initial portion of the string pointed to by *str* to a type **unsigned**  
45277 **long** and **unsigned long long** representation, respectively. First, they decompose the input  
45278 string into three parts:

- 45279 1. An initial, possibly empty, sequence of white-space characters (as specified by *isspace()*)
- 45280 2. A subject sequence interpreted as an integer represented in some radix determined by the  
45281 value of *base*
- 45282 3. A final string of one or more unrecognized characters, including the terminating null byte  
45283 of the input string

45284 Then they shall attempt to convert the subject sequence to an unsigned integer, and return the  
45285 result.45286 If the value of *base* is 0, the expected form of the subject sequence is that of a decimal constant,  
45287 octal constant, or hexadecimal constant, any of which may be preceded by a '+' or '-' sign. A  
45288 decimal constant begins with a non-zero digit, and consists of a sequence of decimal digits. An  
45289 octal constant consists of the prefix '0' optionally followed by a sequence of the digits '0' to  
45290 '7' only. A hexadecimal constant consists of the prefix 0x or 0X followed by a sequence of the  
45291 decimal digits and letters 'a' (or 'A') to 'f' (or 'F') with values 10 to 15 respectively.45292 If the value of *base* is between 2 and 36, the expected form of the subject sequence is a sequence  
45293 of letters and digits representing an integer with the radix specified by *base*, optionally preceded  
45294 by a '+' or '-' sign. The letters from 'a' (or 'A') to 'z' (or 'Z') inclusive are ascribed the  
45295 values 10 to 35; only letters whose ascribed values are less than that of *base* are permitted. If the  
45296 value of *base* is 16, the characters 0x or 0X may optionally precede the sequence of letters and  
45297 digits, following the sign if present.45298 The subject sequence is defined as the longest initial subsequence of the input string, starting  
45299 with the first non-white-space character that is of the expected form. The subject sequence shall  
45300 contain no characters if the input string is empty or consists entirely of white-space characters,  
45301 or if the first non-white-space character is other than a sign or a permissible letter or digit.45302 If the subject sequence has the expected form and the value of *base* is 0, the sequence of  
45303 characters starting with the first digit shall be interpreted as an integer constant. If the subject  
45304 sequence has the expected form and the value of *base* is between 2 and 36, it shall be used as the  
45305 base for conversion, ascribing to each letter its value as given above. If the subject sequence  
45306 begins with a minus sign, the value resulting from the conversion shall be negated. A pointer to  
45307 the final string shall be stored in the object pointed to by *endptr*, provided that *endptr* is not a null  
45308 pointer.

45309 CX In other than the C or POSIX locales, other implementation-defined subject sequences may be  
45310 accepted.

45311 If the subject sequence is empty or does not have the expected form, no conversion shall be |  
45312 performed; the value of *str* shall be stored in the object pointed to by *endptr*, provided that *endptr* |  
45313 is not a null pointer.

45314 CX The *strtoul()* function shall not change the setting of *errno* if successful.

45315 Since 0, {ULONG\_MAX}, and {ULLONG\_MAX} are returned on error and are also valid returns |  
45316 on success, an application wishing to check for error situations should set *errno* to 0, then call |  
45317 *strtoul()* or *strtoull()*, then check *errno*.

#### 45318 RETURN VALUE

45319 Upon successful completion, these functions shall return the converted value, if any. If no  
45320 CX conversion could be performed, 0 shall be returned and *errno* may be set to [EINVAL]. If the  
45321 correct value is outside the range of representable values, {ULONG\_MAX} or {ULLONG\_MAX}  
45322 shall be returned and *errno* set to [ERANGE].

#### 45323 ERRORS

45324 These functions shall fail if:

45325 CX [EINVAL] The value of *base* is not supported.

45326 [ERANGE] The value to be returned is not representable.

45327 These functions may fail if:

45328 CX [EINVAL] No conversion could be performed.

#### 45329 EXAMPLES

45330 None.

#### 45331 APPLICATION USAGE

45332 None.

#### 45333 RATIONALE

45334 None.

#### 45335 FUTURE DIRECTIONS

45336 None.

#### 45337 SEE ALSO

45338 *isalpha()*, *scanf()*, *strtod()*, *strtol()*, the Base Definitions volume of IEEE Std 1003.1-200x,  
45339 <stdlib.h>

#### 45340 CHANGE HISTORY

45341 First released in Issue 4. Derived from the ANSI C standard.

#### 45342 Issue 5

45343 The DESCRIPTION is updated to indicate that *errno* is not changed if the function is successful.

#### 45344 Issue 6

45345 Extensions beyond the ISO C standard are now marked.

45346 The following new requirements on POSIX implementations derive from alignment with the  
45347 Single UNIX Specification:

- 45348 • The [EINVAL] error condition is added for when the value of *base* is not supported.

45349 In the RETURN VALUE and ERRORS sections, the [EINVAL] optional error condition is  
45350 added if no conversion could be performed.

45351 The following changes are made for alignment with the ISO/IEC 9899:1999 standard:

45352 • The *strtoul()* prototype is updated.

45353 • The *strtoull()* function is added.

45354 **NAME**

45355            strtoumax — convert string to integer type

45356 **SYNOPSIS**

45357            #include &lt;inttypes.h&gt;

45358            uintmax\_t strtoumax(const char \*restrict *nptr*, char \*\*restrict *endptr*,  
45359                                int *base*);45360 **DESCRIPTION**45361            Refer to *strtoimax()*.

45362 **NAME**

45363 strxfrm — string transformation

45364 **SYNOPSIS**

45365 #include &lt;string.h&gt;

45366 size\_t strxfrm(char \*restrict *s1*, const char \*restrict *s2*, size\_t *n*);45367 **DESCRIPTION**

45368 CX The functionality described on this reference page is aligned with the ISO C standard. Any  
 45369 conflict between the requirements described here and the ISO C standard is unintentional. This  
 45370 volume of IEEE Std 1003.1-200x defers to the ISO C standard.

45371 The *strxfrm()* function shall transform the string pointed to by *s2* and place the resulting string  
 45372 into the array pointed to by *s1*. The transformation is such that if *strcmp()* is applied to two  
 45373 transformed strings, it shall return a value greater than, equal to, or less than 0, corresponding to  
 45374 the result of *strcoll()* applied to the same two original strings. No more than *n* bytes are placed  
 45375 into the resulting array pointed to by *s1*, including the terminating null byte. If *n* is 0, *s1*  
 45376 is permitted to be a null pointer. If copying takes place between objects that overlap, the behavior  
 45377 is undefined.

45378 CX The *strxfrm()* function shall not change the setting of *errno* if successful.

45379 Since no return value is reserved to indicate an error, an application wishing to check for error  
 45380 situations should set *errno* to 0, then call *strxfrm()*, then check *errno*.

45381 **RETURN VALUE**

45382 Upon successful completion, *strxfrm()* shall return the length of the transformed string (not  
 45383 including the terminating null byte). If the value returned is *n* or more, the contents of the array  
 45384 pointed to by *s1* are unspecified.

45385 CX On error, *strxfrm()* may set *errno* but no return value is reserved to indicate an error.

45386 **ERRORS**

45387 The *strxfrm()* function may fail if:

45388 CX [EINVAL] The string pointed to by the *s2* argument contains characters outside the  
 45389 domain of the collating sequence.

45390 **EXAMPLES**

45391 None.

45392 **APPLICATION USAGE**

45393 The transformation function is such that two transformed strings can be ordered by *strcmp()* as  
 45394 appropriate to collating sequence information in the program's locale (category *LC\_COLLATE*).

45395 The fact that when *n* is 0 *s1* is permitted to be a null pointer is useful to determine the size of the  
 45396 *s1* array prior to making the transformation.

45397 **RATIONALE**

45398 None.

45399 **FUTURE DIRECTIONS**

45400 None.

45401 **SEE ALSO**

45402 *strcmp()*, *strcoll()*, the Base Definitions volume of IEEE Std 1003.1-200x, <string.h>

45403 **CHANGE HISTORY**

45404 First released in Issue 3.

45405 Entry included for alignment with the ISO C standard.

45406 **Issue 5**45407 The DESCRIPTION is updated to indicate that *errno* does not change if the function is successful.  
4540845409 **Issue 6**

45410 Extensions beyond the ISO C standard are now marked.

45411 The following new requirements on POSIX implementations derive from alignment with the  
45412 Single UNIX Specification:

- 45413
- In the RETURN VALUE and ERRORS sections, the [EINVAL] optional error condition is  
45414 added if no conversion could be performed.

45415 The *strxfrm()* prototype is updated for alignment with the ISO/IEC 9899:1999 standard.

45416 **NAME**

45417 swab — swap bytes

45418 **SYNOPSIS**

45419 XSI #include &lt;unistd.h&gt;

45420 void swab(const void \*restrict *src*, void \*restrict *dest*,  
45421 ssize\_t *nbytes*);

45422

45423 **DESCRIPTION**

45424 The *swab()* function shall copy *nbytes* bytes, which are pointed to by *src*, to the object pointed to  
45425 by *dest*, exchanging adjacent bytes. The *nbytes* argument should be even. If *nbytes* is odd, *swab()*  
45426 copies and exchanges *nbytes*–1 bytes and the disposition of the last byte is unspecified. If  
45427 copying takes place between objects that overlap, the behavior is undefined. If *nbytes* is  
45428 negative, *swab()* does nothing.

45429 **RETURN VALUE**

45430 None.

45431 **ERRORS**

45432 No errors are defined.

45433 **EXAMPLES**

45434 None.

45435 **APPLICATION USAGE**

45436 None.

45437 **RATIONALE**

45438 None.

45439 **FUTURE DIRECTIONS**

45440 None.

45441 **SEE ALSO**

45442 The Base Definitions volume of IEEE Std 1003.1-200x, &lt;unistd.h&gt;

45443 **CHANGE HISTORY**

45444 First released in Issue 1. Derived from Issue 1 of the SVID.

45445 **Issue 6**

45446 The **restrict** keyword is added to the *swab()* prototype for alignment with the  
45447 ISO/IEC 9899:1999 standard.



45448 **NAME**

45449 swapcontext — swap user context

45450 **SYNOPSIS**

45451 xSI #include &lt;ucontext.h&gt;

45452 int swapcontext(ucontext\_t \*restrict oucp,  
45453 const ucontext\_t \*restrict ucp);

45454

45455 **DESCRIPTION**45456 Refer to *makecontext()*.

45457 **NAME**

45458           swprintf — print formatted wide-character output

45459 **SYNOPSIS**

45460           #include <stdio.h>

45461           #include <wchar.h>

45462           int swprintf(wchar\_t \*ws, size\_t n, const wchar\_t \*format, ...);

45463 **DESCRIPTION**

45464           Refer to *fwprintf()*.

45465 **NAME**

45466 swscanf — convert formatted wide-character input

45467 **SYNOPSIS**

45468 #include &lt;stdio.h&gt;

45469 #include &lt;wchar.h&gt;

45470 int swscanf(const wchar\_t \*restrict ws, |  
45471 const wchar\_t \*restrict format, ... ); |45472 **DESCRIPTION** |45473 Refer to *fwscanf()*.

45474 **NAME**45475 `symlink` — make symbolic link to a file45476 **SYNOPSIS**45477 `#include <unistd.h>`45478 `int symlink(const char *path1, const char *path2);`45479 **DESCRIPTION**

45480 The `symlink()` function shall create a symbolic link called `path2` that contains the string pointed  
 45481 to by `path1` (`path2` is the name of the symbolic link created, `path1` is the string contained in the  
 45482 symbolic link).

45483 The string pointed to by `path1` shall be treated only as a character string and shall not be  
 45484 validated as a pathname.

45485 If the `symlink()` function fails for any reason other than [EIO], any file named by `path2` shall be  
 45486 unaffected.

45487 **RETURN VALUE**

45488 Upon successful completion, `symlink()` shall return 0; otherwise, it shall return `-1` and set `errno` to  
 45489 indicate the error.

45490 **ERRORS**45491 The `symlink()` function shall fail if:

45492 [EACCES] Write permission is denied in the directory where the symbolic link is being  
 45493 created, or search permission is denied for a component of the path prefix of  
 45494 `path2`.

45495 [EEXIST] The `path2` argument names an existing file or symbolic link.

45496 [EIO] An I/O error occurs while reading from or writing to the file system.

45497 [ELOOP] A loop exists in symbolic links encountered during resolution of the `path2`  
 45498 argument.

45499 [ENAMETOOLONG]

45500 The length of the `path2` argument exceeds {PATH\_MAX} or a pathname  
 45501 component is longer than {NAME\_MAX} or the length of the `path1` argument  
 45502 is longer than {SYMLINK\_MAX}.

45503 [ENOENT] A component of `path2` does not name an existing file or `path2` is an empty  
 45504 string.

45505 [ENOSPC] The directory in which the entry for the new symbolic link is being placed  
 45506 cannot be extended because no space is left on the file system containing the  
 45507 directory, or the new symbolic link cannot be created because no space is left  
 45508 on the file system which shall contain the link, or the file system is out of file-  
 45509 allocation resources.

45510 [ENOTDIR] A component of the path prefix of `path2` is not a directory.

45511 [EROFS] The new symbolic link would reside on a read-only file system.

45512 The `symlink()` function may fail if:

45513 [ELOOP] More than {SYMLOOP\_MAX} symbolic links were encountered during  
 45514 resolution of the `path2` argument.

45515 [ENAMETOOLONG]

45516 As a result of encountering a symbolic link in resolution of the `path2`

45517 argument, the length of the substituted pathname string exceeded |  
45518 {PATH\_MAX} bytes (including the terminating null byte), or the length of the  
45519 string pointed to by *path1* exceeded {SYMLINK\_MAX}.

**45520 EXAMPLES**

45521 None.

**45522 APPLICATION USAGE**

45523 Like a hard link, a symbolic link allows a file to have multiple logical names. The presence of a  
45524 hard link guarantees the existence of a file, even after the original name has been removed. A  
45525 symbolic link provides no such assurance; in fact, the file named by the *path1* argument need not  
45526 exist when the link is created. A symbolic link can cross file system boundaries.

45527 Normal permission checks are made on each component of the symbolic link pathname during |  
45528 its resolution. |

**45529 RATIONALE**

45530 Since IEEE Std 1003.1-200x does not require any association of file times with symbolic links,  
45531 there is no requirement that file times be updated by *symlink()*.

**45532 FUTURE DIRECTIONS**

45533 None.

**45534 SEE ALSO**

45535 *lchown()*, *link()*, *lstat()*, *open()*, *readlink()*, *unlink()*, the Base Definitions volume of  
45536 IEEE Std 1003.1-200x, <**unistd.h**>

**45537 CHANGE HISTORY**

45538 First released in Issue 4, Version 2.

**45539 Issue 5**

45540 Moved from X/OPEN UNIX extension to BASE.

**45541 Issue 6**

45542 The following changes were made to align with the IEEE P1003.1a draft standard: |

- 45543
- The DESCRIPTION text is updated.
  - The [ELOOP] optional error condition is added.
- 45544

45545 **NAME**

45546 sync — schedule file system updates

45547 **SYNOPSIS**

45548 XSI #include &lt;unistd.h&gt;

45549 void sync(void);

45550

45551 **DESCRIPTION**45552 The *sync()* function shall cause all information in memory that updates file systems to be  
45553 scheduled for writing out to all file systems.45554 The writing, although scheduled, is not necessarily complete upon return from *sync()*.45555 **RETURN VALUE**45556 The *sync()* function shall not return a value.45557 **ERRORS**

45558 No errors are defined.

45559 **EXAMPLES**

45560 None.

45561 **APPLICATION USAGE**

45562 None.

45563 **RATIONALE**

45564 None.

45565 **FUTURE DIRECTIONS**

45566 None.

45567 **SEE ALSO**45568 *fsync()*, the Base Definitions volume of IEEE Std 1003.1-200x, <unistd.h>45569 **CHANGE HISTORY**

45570 First released in Issue 4, Version 2.

45571 **Issue 5**

45572 Moved from X/OPEN UNIX extension to BASE.

45573 **NAME**

45574 sysconf — get configurable system variables

45575 **SYNOPSIS**

45576 #include &lt;unistd.h&gt;

45577 long sysconf(int name);

45578 **DESCRIPTION**

45579 The *sysconf()* function provides a method for the application to determine the current value of a  
 45580 configurable system limit or option (*variable*). Support for some system variables is dependent  
 45581 on implementation options (as indicated by the margin codes in the following table). Where an  
 45582 implementation option is not supported, the variable need not be supported.

45583 The *name* argument represents the system variable to be queried. The following table lists the  
 45584 minimal set of system variables from <limits.h> or <unistd.h> that can be returned by *sysconf()*,  
 45585 and the symbolic constants, defined in <unistd.h> that are the corresponding values used for  
 45586 *name*. Support for some configuration variables is dependent on implementation options (see  
 45587 shading and margin codes in the table below). Where an implementation option is not  
 45588 supported, the variable need not be supported.

45589

45590

45591 AIO

45592

45593

45594

45595 XSI

45596

45597

45598

45599

45600

45601

45602

45603 XSI

45604

45605

45606 XSI

45607

45608

45609

45610 TSF

45611

45612

45613

45614 MSG

45615

45616

	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
	{HOST_NAME_MAX}	_SC_HOST_NAME_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 <i>getgrnam_r()</i> data buffers	_SC_GETGR_R_SIZE_MAX
	Maximum size of <i>getpwuid_r()</i> and <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

	Variable	Value of Name
45617		
45618		
45619 ADV	_POSIX_ADVISORY_INFO	_SC_ADVISORY_INFO
45620 BAR	_POSIX_BARRIERS	_SC_BARRIERS
45621 AIO	_POSIX_ASYNCHRONOUS_IO	_SC_ASYNCHRONOUS_IO
45622	_POSIX_BASE	_SC_BASE
45623	_POSIX_C_LANG_SUPPORT	_SC_C_LANG_SUPPORT
45624	_POSIX_C_LANG_SUPPORT_R	_SC_C_LANG_SUPPORT_R
45625 CS	_POSIX_CLOCK_SELECTION	_SC_CLOCK_SELECTION
45626 CPT	_POSIX_CPUTIME	_SC_CPUTIME
45627	_POSIX_DEVICE_IO	_SC_DEVICE_IO
45628	_POSIX_DEVICE_SPECIFIC	_SC_DEVICE_SPECIFIC
45629	_POSIX_DEVICE_SPECIFIC_R	_SC_DEVICE_SPECIFIC_R
45630	_POSIX_FD_MGMT	_SC_FD_MGMT
45631	_POSIX_FIFO	_SC_FIFO
45632	_POSIX_FILE_ATTRIBUTES	_SC_FILE_ATTRIBUTES
45633	_POSIX_FILE_LOCKING	_SC_FILE_LOCKING
45634	_POSIX_FILE_SYSTEM	_SC_FILE_SYSTEM
45635 FSC	_POSIX_FSYNC	_SC_FSYNC
45636	_POSIX_JOB_CONTROL	_SC_JOB_CONTROL
45637 MF	_POSIX_MAPPED_FILES	_SC_MAPPED_FILES
45638 ML	_POSIX_MEMLOCK	_SC_MEMLOCK
45639 MLR	_POSIX_MEMLOCK_RANGE	_SC_MEMLOCK_RANGE
45640 MPR	_POSIX_MEMORY_PROTECTION	_SC_MEMORY_PROTECTION
45641 MSG	_POSIX_MESSAGE_PASSING	_SC_MESSAGE_PASSING
45642 MON	_POSIX_MONOTONIC_CLOCK	_SC_MONOTONIC_CLOCK
45643	_POSIX_MULTI_PROCESS	_SC_MULTI_PROCESS
45644	_POSIX_NETWORKING	_SC_NETWORKING
45645	_POSIX_PIPE	_SC_PIPE
45646 PIO	_POSIX_PRIORITIZED_IO	_SC_PRIORITIZED_IO
45647 PS	_POSIX_PRIORITY_SCHEDULING	_SC_PRIORITY_SCHEDULING
45648 THR	_POSIX_READER_WRITER_LOCKS	_SC_READER_WRITER_LOCKS
45649 RTS	_POSIX_REALTIME_SIGNALS	_SC_REALTIME_SIGNALS
45650	_POSIX_REGEX	_SC_REGEX
45651	_POSIX_SAVED_IDS	_SC_SAVED_IDS
45652 SEM	_POSIX_SEMAPHORES	_SC_SEMAPHORES
45653 SHM	_POSIX_SHARED_MEMORY_OBJECTS	_SC_SHARED_MEMORY_OBJECTS
45654	_POSIX_SHELL	_SC_SHELL
45655	_POSIX_SIGNALS	_SC_SIGNALS
45656	_POSIX_SINGLE_PROCESS	_SC_SINGLE_PROCESS
45657 SPN	_POSIX_SPAWN	_SC_SPAWN
45658 SPI	_POSIX_SPIN_LOCKS	_SC_SPIN_LOCKS
45659 SS	_POSIX_SPORADIC_SERVER	_SC_SPORADIC_SERVER
45660 SIO	_POSIX_SYNCHRONIZED_IO	_SC_SYNCHRONIZED_IO
45661	_POSIX_SYSTEM_DATABASE	_SC_SYSTEM_DATABASE
45662	_POSIX_SYSTEM_DATABASE_R	_SC_SYSTEM_DATABASE_R



45663

45664

45665 TSA

45666 TSS

45667 TCT

45668 TPI

45669 TPP

45670 TPS

45671 TSH

45672 TSF

45673 TSP

45674 THR

45675 TMO

45676 TMR

45677 TRC

45678 TEF

45679 TRI

45680 TRL

45681 TYM

45682

45683

45684

45685

45686

45687

45688

45689

45690

45691

45692

45693

45694

45695

45696 BE

45697

45698

45699

45700

45701

45702

45703

45704

45705 XSI

45706

	Variable	Value of Name
	_POSIX_THREAD_ATTR_STACKADDR	_SC_THREAD_ATTR_STACKADDR
	_POSIX_THREAD_ATTR_STACKSIZE	_SC_THREAD_ATTR_STACKSIZE
	_POSIX_THREAD_CPU_TIME	_SC_THREAD_CPU_TIME
	_POSIX_THREAD_PRIO_INHERIT	_SC_THREAD_PRIO_INHERIT
	_POSIX_THREAD_PRIO_PROTECT	_SC_THREAD_PRIO_PROTECT
	_POSIX_THREAD_PRIORITY_SCHEDULING	_SC_THREAD_PRIORITY_SCHEDULING
	_POSIX_THREAD_PROCESS_SHARED	_SC_THREAD_PROCESS_SHARED
	_POSIX_THREAD_SAFE_FUNCTIONS	_SC_THREAD_SAFE_FUNCTIONS
	_POSIX_THREAD_SPORADIC_SERVER	_SC_THREAD_SPORADIC_SERVER
	_POSIX_THREADS	_SC_THREADS
	_POSIX_TIMEOUTS	_SC_TIMEOUTS
	_POSIX_TIMERS	_SC_TIMERS
	_POSIX_TRACE	_SC_TRACE
	_POSIX_TRACE_EVENT_FILTER	_SC_TRACE_EVENT_FILTER
	_POSIX_TRACE_INHERIT	_SC_TRACE_INHERIT
	_POSIX_TRACE_LOG	_SC_TRACE_LOG
	_POSIX_TYPED_MEMORY_OBJECTS	_SC_TYPED_MEMORY_OBJECTS
	_POSIX_USER_GROUPS	_SC_USER_GROUPS
	_POSIX_USER_GROUPS_R	_SC_USER_GROUPS_R
	_POSIX_VERSION	_SC_VERSION
	_POSIX_V6_ILP32_OFF32	_SC_V6_ILP32_OFF32
	_POSIX_V6_ILP32_OFFBIG	_SC_V6_ILP32_OFFBIG
	_POSIX_V6_LP64_OFF64	_SC_V6_LP64_OFF64
	_POSIX_V6_LPBIG_OFFBIG	_SC_V6_LPBIG_OFFBIG
	_POSIX2_C_BIND	_SC_2_C_BIND
	_POSIX2_C_DEV	_SC_2_C_DEV
	_POSIX2_C_VERSION	_SC_2_C_VERSION
	_POSIX2_CHAR_TERM	_SC_2_CHAR_TERM
	_POSIX2_FORT_DEV	_SC_2_FORT_DEV
	_POSIX2_FORT_RUN	_SC_2_FORT_RUN
	_POSIX2_LOCALEDEF	_SC_2_LOCALEDEF
	_POSIX2_PBS	_SC_2_PBS
	_POSIX2_PBS_ACCOUNTING	_SC_2_PBS_ACCOUNTING
	_POSIX2_PBS_LOCATE	_SC_2_PBS_LOCATE
	_POSIX2_PBS_MESSAGE	_SC_2_PBS_MESSAGE
	_POSIX2_PBS_TRACK	_SC_2_PBS_TRACK
	_POSIX2_SW_DEV	_SC_2_SW_DEV
	_POSIX2_UPE	_SC_2_UPE
	_POSIX2_VERSION	_SC_2_VERSION
	_REGEX_VERSION	_SC_REGEX_VERSION
	{PAGE_SIZE}	_SC_PAGE_SIZE
	{PAGESIZE}	_SC_PAGESIZE

	Variable	Value of Name
45707		
45708		
45709 THR	{PTHREAD_DESTRUCTOR_ITERATIONS}	_SC_THREAD_DESTRUCTOR_ITERATIONS
45710	{PTHREAD_KEYS_MAX}	_SC_THREAD_KEYS_MAX
45711	{PTHREAD_STACK_MIN}	_SC_THREAD_STACK_MIN
45712	{PTHREAD_THREADS_MAX}	_SC_THREAD_THREADS_MAX
45713	{RE_DUP_MAX}	_SC_RE_DUP_MAX
45714 RTS	{RTSIG_MAX}	_SC_RTSIG_MAX
45715 SEM	{SEM_NSEMS_MAX}	_SC_SEM_NSEMS_MAX
45716	{SEM_VALUE_MAX}	_SC_SEM_VALUE_MAX
45717 RTS	{SIGQUEUE_MAX}	_SC_SIGQUEUE_MAX
45718	{STREAM_MAX}	_SC_STREAM_MAX
45719	{SYMLOOP_MAX}	_SC_SYMLOOP_MAX
45720 TMR	{TIMER_MAX}	_SC_TIMER_MAX
45721	{TTY_NAME_MAX}	_SC_TTY_NAME_MAX
45722	{TZNAME_MAX}	_SC_TZNAME_MAX
45723 XSI	_XBS5_ILP32_OFF32 (LEGACY)	_SC_XBS5_ILP32_OFF32 (LEGACY)
45724	_XBS5_ILP32_OFFBIG (LEGACY)	_SC_XBS5_ILP32_OFFBIG (LEGACY)
45725	_XBS5_LP64_OFF64 (LEGACY)	_SC_XBS5_LP64_OFF64 (LEGACY)
45726	_XBS5_LPBIG_OFFBIG (LEGACY)	_SC_XBS5_LPBIG_OFFBIG (LEGACY)
45727	_XOPEN_CRYPT	_SC_XOPEN_CRYPT
45728	_XOPEN_ENH_I18N	_SC_XOPEN_ENH_I18N
45729	_XOPEN_LEGACY	_SC_XOPEN_LEGACY
45730	_XOPEN_REALTIME	_SC_XOPEN_REALTIME
45731	_XOPEN_REALTIME_THREADS	_SC_XOPEN_REALTIME_THREADS
45732	_XOPEN_SHM	_SC_XOPEN_SHM
45733	_XOPEN_UNIX	_SC_XOPEN_UNIX
45734	_XOPEN_VERSION	_SC_XOPEN_VERSION
45735	_XOPEN_XCU_VERSION	_SC_XOPEN_XCU_VERSION

45736 **RETURN VALUE**

45737 If *name* is an invalid value, *sysconf()* shall return  $-1$  and set *errno* to indicate the error. If the  
 45738 variable corresponding to *name* has no limit, *sysconf()* shall return  $-1$  without changing the value  
 45739 of *errno*. Note that indefinite limits do not imply infinite limits; see <**limits.h**>.

45740 Otherwise, *sysconf()* shall return the current variable value on the system. The value returned  
 45741 shall not be more restrictive than the corresponding value described to the application when it  
 45742 was compiled with the implementation's <**limits.h**> or <**unistd.h**>. The value shall not change  
 45743 during the lifetime of the calling process.

45744 **ERRORS**

45745 The *sysconf()* function shall fail if:

45746 [EINVAL] The value of the *name* argument is invalid.

45747 **EXAMPLES**

45748 None.

45749 **APPLICATION USAGE**

45750 As  $-1$  is a permissible return value in a successful situation, an application wishing to check for  
 45751 error situations should set *errno* to 0, then call *sysconf()*, and, if it returns  $-1$ , check to see if *errno*  
 45752 is non-zero.

45753 If the value of *sysconf(\_SC\_2\_VERSION)* is not equal to the value of the *\_POSIX2\_VERSION*  
 45754 symbolic constant, the utilities available via *system()* or *popen()* might not behave as described in  
 45755 the Shell and Utilities volume of IEEE Std 1003.1-200x. This would mean that the application is

45756 not running in an environment that conforms to the Shell and Utilities volume of  
45757 IEEE Std 1003.1-200x. Some applications might be able to deal with this, others might not.  
45758 However, the functions defined in this volume of IEEE Std 1003.1-200x continue to operate as  
45759 specified, even if: *sysconf(SC\_2\_VERSION)* reports that the utilities no longer perform as  
45760 specified.

#### 45761 RATIONALE

45762 This functionality was added in response to requirements of application developers and of  
45763 system vendors who deal with many international system configurations. It is closely related to  
45764 *pathconf()* and *fpathconf()*.

45765 Although a conforming application can run on all systems by never demanding more resources  
45766 than the minimum values published in this volume of IEEE Std 1003.1-200x, it is useful for that  
45767 application to be able to use the actual value for the quantity of a resource available on any  
45768 given system. To do this, the application makes use of the value of a symbolic constant in  
45769 *<limits.h>* or *<unistd.h>*.

45770 However, once compiled, the application must still be able to cope if the amount of resource  
45771 available is increased. To that end, an application may need a means of determining the quantity  
45772 of a resource, or the presence of an option, at execution time.

45773 Two examples are offered:

- 45774 1. Applications may wish to act differently on systems with or without job control.  
45775 Applications vendors who wish to distribute only a single binary package to all instances  
45776 of a computer architecture would be forced to assume job control is never available if it  
45777 were to rely solely on the *<unistd.h>* value published in this volume of  
45778 IEEE Std 1003.1-200x.
- 45779 2. International applications vendors occasionally require knowledge of the number of clock  
45780 ticks per second. Without these facilities, they would be required to either distribute their  
45781 applications partially in source form or to have 50Hz and 60Hz versions for the various  
45782 countries in which they operate.

45783 It is the knowledge that many applications are actually distributed widely in executable form  
45784 that leads to this facility. If limited to the most restrictive values in the headers, such  
45785 applications would have to be prepared to accept the most limited environments offered by the  
45786 smallest microcomputers. Although this is entirely portable, there was a consensus that they  
45787 should be able to take advantage of the facilities offered by large systems, without the  
45788 restrictions associated with source and object distributions.

45789 During the discussions of this feature, it was pointed out that it is almost always possible for an  
45790 application to discern what a value might be at runtime by suitably testing the various functions  
45791 themselves. And, in any event, it could always be written to adequately deal with error returns  
45792 from the various functions. In the end, it was felt that this imposed an unreasonable level of  
45793 complication and sophistication on the application writer.

45794 This runtime facility is not meant to provide ever-changing values that applications have to  
45795 check multiple times. The values are seen as changing no more frequently than once per system  
45796 initialization, such as by a system administrator or operator with an automatic configuration  
45797 program. This volume of IEEE Std 1003.1-200x specifies that they shall not change within the  
45798 lifetime of the process.

45799 Some values apply to the system overall and others vary at the file system or directory level. The  
45800 latter are described in *pathconf()*.

45801 Note that all values returned must be expressible as integers. String values were considered, but  
45802 the additional flexibility of this approach was rejected due to its added complexity of

- 45803 implementation and use.
- 45804 Some values, such as {PATH\_MAX}, are sometimes so large that they must not be used to, say,  
45805 allocate arrays. The *sysconf()* function returns a negative value to show that this symbolic  
45806 constant is not even defined in this case.
- 45807 Similar to *pathconf()*, this permits the implementation not to have a limit. When one resource is  
45808 infinite, returning an error indicating that some other resource limit has been reached is  
45809 conforming behavior.
- 45810 **FUTURE DIRECTIONS**
- 45811 None.
- 45812 **SEE ALSO**
- 45813 *confstr()*, *pathconf()*, the Base Definitions volume of IEEE Std 1003.1-200x, **<limits.h>**,  
45814 **<unistd.h>**, the Shell and Utilities volume of IEEE Std 1003.1-200x, *getconf*
- 45815 **CHANGE HISTORY**
- 45816 First released in Issue 3.
- 45817 Entry included for alignment with the POSIX.1-1988 standard.
- 45818 **Issue 5**
- 45819 The DESCRIPTION is updated for alignment with the POSIX Realtime Extension and the POSIX  
45820 Threads Extension.
- 45821 The **\_XBS\_** variables and name values are added to the table of system variables in the  
45822 DESCRIPTION. These are all marked EX.
- 45823 **Issue 6**
- 45824 The symbol CLK\_TCK is obsolescent and removed. It is replaced with the phrase “clock ticks  
45825 per second”.
- 45826 The symbol {PASS\_MAX} is removed.
- 45827 The following changes were made to align with the IEEE P1003.1a draft standard:
- 45828 • Table entries added for the following variables: **\_SC\_REGEX**, **\_SC\_SHELL**,  
45829 **\_SC\_REGEX\_VERSION**, **\_SC\_SYMLoop\_MAX**.
- 45830 The following *sysconf()* variables and their associated names are added for alignment with  
45831 IEEE Std 1003.1d-1999:
- 45832 **\_POSIX\_ADVISORY\_INFO**  
45833 **\_POSIX\_CPUTIME**  
45834 **\_POSIX\_SPAWN**  
45835 **\_POSIX\_SPORADIC\_SERVER**  
45836 **\_POSIX\_THREAD\_CPUTIME**  
45837 **\_POSIX\_THREAD\_SPORADIC\_SERVER**  
45838 **\_POSIX\_TIMEOUTS**
- 45839 The following changes are made to the DESCRIPTION for alignment with IEEE Std 1003.1j-2000:
- 45840 • A statement expressing the dependency of support for some system variables on  
45841 implementation options is added.
- 45842 • The following system variables are added:

45843            \_POSIX\_BARRIERS  
45844            \_POSIX\_CLOCK\_SELECTION  
45845            \_POSIX\_MONOTONIC\_CLOCK  
45846            \_POSIX\_READER\_WRITER\_LOCKS  
45847            \_POSIX\_SPIN\_LOCKS  
45848            \_POSIX\_TYPED\_MEMORY\_OBJECTS

45849            The following system variables are added for alignment with IEEE Std 1003.2d-1994:

45850            \_POSIX2\_PBS  
45851            \_POSIX2\_PBS\_ACCOUNTING  
45852            \_POSIX2\_PBS\_LOCATE  
45853            \_POSIX2\_PBS\_MESSAGE  
45854            \_POSIX2\_PBS\_TRACK

45855            The following *sysconf()* variables and their associated names are added for alignment with  
45856            IEEE Std 1003.1q-2000:

45857            \_POSIX\_TRACE  
45858            \_POSIX\_TRACE\_EVENT\_FILTER  
45859            \_POSIX\_TRACE\_INHERIT  
45860            \_POSIX\_TRACE\_LOG

45861            The macros associated with the *c89* programming models are marked LEGACY, and new  
45862            equivalent macros associated with *c99* are introduced.

45863 **NAME**

45864            syslog — log a message

45865 **SYNOPSIS**

45866 XSI        #include <syslog.h>

45867            void syslog(int *priority*, const char *\*message*, ... /\* *argument* \*/);

45868

45869 **DESCRIPTION**

45870            Refer to *closelog*().

45871 **NAME**

45872           system — issue a command

45873 **SYNOPSIS**

45874           #include &lt;stdlib.h&gt;

45875           int system(const char \**command*);45876 **DESCRIPTION**

45877 CX       The functionality described on this reference page is aligned with the ISO C standard. Any  
 45878           conflict between the requirements described here and the ISO C standard is unintentional. This  
 45879           volume of IEEE Std 1003.1-200x defers to the ISO C standard.

45880           If *command* is a null pointer, the *system()* function shall determine whether the host environment  
 45881           has a command processor. If *command* is not a null pointer, the *system()* function shall pass the  
 45882           string pointed to by *command* to that command processor to be executed in an implementation-  
 45883           defined manner; this might then cause the program calling *system()* to behave in a non-  
 45884           conforming manner or to terminate.

45885 CX       The environment of the executed command shall be as if a child process were created using  
 45886           *fork()*, and the child process invoked the *sh* utility using *execl()* as follows:

```
45887           execl(<shell path>, "sh", "-c", command, (char *)0);
```

45888           where <shell path> is an unspecified pathname for the *sh* utility.

45889           The *system()* function shall ignore the SIGINT and SIGQUIT signals, and shall block the  
 45890           SIGCHLD signal, while waiting for the command to terminate. If this might cause the  
 45891           application to miss a signal that would have killed it, then the application should examine the  
 45892           return value from *system()* and take whatever action is appropriate to the application if the  
 45893           command terminated due to receipt of a signal.

45894           The *system()* function shall not affect the termination status of any child of the calling processes  
 45895           other than the process or processes it itself creates.

45896           The *system()* function shall not return until the child process has terminated.

45897 **RETURN VALUE**

45898           If *command* is a null pointer, *system()* shall return non-zero to indicate that a command processor  
 45899 CX           is available, or zero if none is available. The *system()* function shall always return non-zero when  
 45900           *command* is NULL.

45901 CX       If *command* is not a null pointer, *system()* shall return the termination status of the command  
 45902           language interpreter in the format specified by *waitpid()*. The termination status shall be as  
 45903           defined for the *sh* utility; otherwise, the termination status is unspecified. If some error prevents  
 45904           the command language interpreter from executing after the child process is created, the return  
 45905           value from *system()* shall be as if the command language interpreter had terminated using  
 45906           *exit(127)* or *\_exit(127)*. If a child process cannot be created, or if the termination status for the  
 45907           command language interpreter cannot be obtained, *system()* shall return -1 and set *errno* to  
 45908           indicate the error.

45909 **ERRORS**

45910 CX       The *system()* function may set *errno* values as described by *fork()*.

45911           In addition, *system()* may fail if:

45912 CX       [ECHILD]       The status of the child process created by *system()* is no longer available.

45913 **EXAMPLES**

45914 None.

45915 **APPLICATION USAGE**

45916 If the return value of *system()* is not `-1`, its value can be decoded through the use of the macros  
45917 described in `<sys/wait.h>`. For convenience, these macros are also provided in `<stdlib.h>`.

45918 Note that, while *system()* must ignore `SIGINT` and `SIGQUIT` and block `SIGCHLD` while waiting  
45919 for the child to terminate, the handling of signals in the executed command is as specified by  
45920 *fork()* and *exec*. For example, if `SIGINT` is being caught or is set to `SIG_DFL` when *system()* is  
45921 called, then the child is started with `SIGINT` handling set to `SIG_DFL`.

45922 Ignoring `SIGINT` and `SIGQUIT` in the parent process prevents coordination problems (two  
45923 processes reading from the same terminal, for example) when the executed command ignores or  
45924 catches one of the signals. It is also usually the correct action when the user has given a  
45925 command to the application to be executed synchronously (as in the `'!'` command in many  
45926 interactive applications). In either case, the signal should be delivered only to the child process,  
45927 not to the application itself. There is one situation where ignoring the signals might have less  
45928 than the desired effect. This is when the application uses *system()* to perform some task invisible  
45929 to the user. If the user typed the interrupt character ("`^C`", for example) while *system()* is being  
45930 used in this way, one would expect the application to be killed, but only the executed command  
45931 is killed. Applications that use *system()* in this way should carefully check the return status from  
45932 *system()* to see if the executed command was successful, and should take appropriate action  
45933 when the command fails.

45934 Blocking `SIGCHLD` while waiting for the child to terminate prevents the application from  
45935 catching the signal and obtaining status from *system()*'s child process before *system()* can get the  
45936 status itself.

45937 The context in which the utility is ultimately executed may differ from that in which *system()*  
45938 was called. For example, file descriptors that have the `FD_CLOEXEC` flag set are closed, and the  
45939 process ID and parent process ID are different. Also, if the executed utility changes its  
45940 environment variables or its current working directory, that change is not reflected in the caller's  
45941 context.

45942 There is no defined way for an application to find the specific path for the shell. However,  
45943 *confstr()* can provide a value for *PATH* that is guaranteed to find the *sh* utility.

45944 **RATIONALE**

45945 The *system()* function should not be used by programs that have set user (or group) ID  
45946 privileges. The *fork()* and *exec* family of functions (except *execlp()* and *execvp()*), should be used  
45947 instead. This prevents any unforeseen manipulation of the environment of the user that could  
45948 cause execution of commands not anticipated by the calling program.

45949 There are three levels of specification for the *system()* function. The ISO C standard gives the  
45950 most basic. It requires that the function exists, and defines a way for an application to query  
45951 whether a command language interpreter exists. It says nothing about the command language or  
45952 the environment in which the command is interpreted.

45953 IEEE Std 1003.1-200x places additional restrictions on *system()*. It requires that if there is a  
45954 command language interpreter, the environment must be as specified by *fork()* and *exec*. This  
45955 ensures, for example, that close-on-exec works, that file locks are not inherited, and that the  
45956 process ID is different. It also specifies the return value from *system()* when the command line  
45957 can be run, thus giving the application some information about the command's completion  
45958 statu.



45959 Finally, IEEE Std 1003.1-200x requires the command to be interpreted as in the shell command  
45960 language defined in the Shell and Utilities volume of IEEE Std 1003.1-200x.

45961 Note that, *system*(NULL) is required to return non-zero, indicating that there is a command  
45962 language interpreter. At first glance, this would seem to conflict with the ISO C standard which  
45963 allows *system*(NULL) to return zero. There is no conflict, however. A system must have a  
45964 command language interpreter, and is non-conforming if none is present. It is therefore  
45965 permissible for the *system*() function on such a system to implement the behavior specified by  
45966 the ISO C standard as long as it is understood that the implementation does not conform to  
45967 IEEE Std 1003.1-200x if *system*(NULL) returns zero.

45968 It was explicitly decided that when *command* is NULL, *system*() should not be required to check  
45969 to make sure that the command language interpreter actually exists with the correct mode, that  
45970 there are enough processes to execute it, and so on. The call *system*(NULL) could, theoretically,  
45971 check for such problems as too many existing child processes, and return zero. However, it  
45972 would be inappropriate to return zero due to such a (presumably) transient condition. If some  
45973 condition exists that is not under the control of this application and that would cause any  
45974 *system*() call to fail, that system has been rendered non-conforming.

45975 Early drafts required, or allowed, *system*() to return with *errno* set to [EINTR] if it was  
45976 interrupted with a signal. This error return was removed, and a requirement that *system*() not  
45977 return until the child has terminated was added. This means that if a *waitpid*() call in *system*()  
45978 exits with *errno* set to [EINTR], *system*() must re-issue the *waitpid*(). This change was made for  
45979 two reasons:

- 45980 1. There is no way for an application to clean up if *system*() returns [EINTR], short of calling  
45981 *wait*(), and that could have the undesirable effect of returning the status of children other  
45982 than the one started by *system*()).
- 45983 2. While it might require a change in some historical implementations, those  
45984 implementations already have to be changed because they use *wait*() instead of *waitpid*()).

45985 Note that if the application is catching SIGCHLD signals, it will receive such a signal before a  
45986 successful *system*() call returns.

45987 To conform to IEEE Std 1003.1-200x, *system*() must use *waitpid*(), or some similar function,  
45988 instead of *wait*()).

45989 The following code sample illustrates how *system*() might be implemented on an  
45990 implementation conforming to IEEE Std 1003.1-200x.

```
45991 #include <signal.h>
45992 int system(const char *cmd)
45993 {
45994     int stat;
45995     pid_t pid;
45996     struct sigaction sa, savintr, savequit;
45997     sigset_t saveblock;
45998     if (cmd == NULL)
45999         return(1);
46000     sa.sa_handler = SIG_IGN;
46001     sigemptyset(&sa.sa_mask);
46002     sa.sa_flags = 0;
46003     sigemptyset(&savintr.sa_mask);
46004     sigemptyset(&savequit.sa_mask);
46005     sigaction(SIGINT, &sa, &savintr);
46006     sigaction(SIGQUIT, &sa, &savequit);
```

```

46007     sigaddset(&sa.sa_mask, SIGCHLD);
46008     sigprocmask(SIG_BLOCK, &sa.sa_mask, &saveblock);
46009     if ((pid = fork()) == 0) {
46010         sigaction(SIGINT, &saveintr, (struct sigaction *)0);
46011         sigaction(SIGQUIT, &savequit, (struct sigaction *)0);
46012         sigprocmask(SIG_SETMASK, &saveblock, (sigset_t *)0);
46013         execl("/bin/sh", "sh", "-c", cmd, (char *)0);
46014         _exit(127);
46015     }
46016     if (pid == -1) {
46017         stat = -1; /* errno comes from fork() */
46018     } else {
46019         while (waitpid(pid, &stat, 0) == -1) {
46020             if (errno != EINTR){
46021                 stat = -1;
46022                 break;
46023             }
46024         }
46025     }
46026     sigaction(SIGINT, &saveintr, (struct sigaction *)0);
46027     sigaction(SIGQUIT, &savequit, (struct sigaction *)0);
46028     sigprocmask(SIG_SETMASK, &saveblock, (sigset_t *)0);
46029     return(stat);
46030 }

```

46031 Note that, while a particular implementation of *system()* (such as the one above) can assume a  
46032 particular path for the shell, such a path is not necessarily valid on another system. The above  
46033 example is not portable, and is not intended to be.

46034 One reviewer suggested that an implementation of *system()* might want to use an environment  
46035 variable such as *SHELL* to determine which command interpreter to use. The supposed  
46036 implementation would use the default command interpreter if the one specified by the  
46037 environment variable was not available. This would allow a user, when using an application  
46038 that prompts for command lines to be processed using *system()*, to specify a different command  
46039 interpreter. Such an implementation is discouraged. If the alternate command interpreter did not  
46040 follow the command line syntax specified in the Shell and Utilities volume of  
46041 IEEE Std 1003.1-200x, then changing *SHELL* would render *system()* non-conforming. This would  
46042 affect applications that expected the specified behavior from *system()*, and since the Shell and  
46043 Utilities volume of IEEE Std 1003.1-200x does not mention that *SHELL* affects *system()*, the  
46044 application would not know that it needed to unset *SHELL*.

#### 46045 FUTURE DIRECTIONS

46046 None.

#### 46047 SEE ALSO

46048 *exec*, *pipe()*, *waitpid()*, the Base Definitions volume of IEEE Std 1003.1-200x, <**limits.h**>,  
46049 <**signal.h**>, <**stdlib.h**>, <**sys/wait.h**>, the Shell and Utilities volume of IEEE Std 1003.1-200x, *sh* |

#### 46050 CHANGE HISTORY

46051 First released in Issue 1. Derived from Issue 1 of the SVID. |

46052 **Issue 6**

46053

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

46054

- The DESCRIPTION is adjusted to reflect the behavior on systems that do not support the Shell option.

46055

## 46056 NAME

46057 tan, tanf, tanl — tangent function

## 46058 SYNOPSIS

46059 #include &lt;math.h&gt;

46060 double tan(double x);

46061 float tanf(float x);

46062 long double tanl(long double x);

## 46063 DESCRIPTION

46064 CX The functionality described on this reference page is aligned with the ISO C standard. Any  
 46065 conflict between the requirements described here and the ISO C standard is unintentional. This  
 46066 volume of IEEE Std 1003.1-200x defers to the ISO C standard.

46067 These functions shall compute the tangent of their argument *x*, measured in radians.

46068 An application wishing to check for error situations should set *errno* to zero and call  
 46069 *feclearexcept*(FE\_ALL\_EXCEPT) before calling these functions. On return, if *errno* is non-zero or  
 46070 *fetestexcept*(FE\_INVALID | FE\_DIVBYZERO | FE\_OVERFLOW | FE\_UNDERFLOW) is non-  
 46071 zero, an error has occurred.

## 46072 RETURN VALUE

46073 Upon successful completion, these functions shall return the tangent of *x*.

46074 If the correct value would cause underflow, and is not representable, a range error may occur,  
 46075 MX and either 0.0 (if supported), or an implementation-defined value shall be returned.

46076 MX If *x* is NaN, a NaN shall be returned.46077 If *x* is  $\pm 0$ , *x* shall be returned.46078 If *x* is subnormal, a range error may occur and *x* should be returned.

46079 If *x* is  $\pm\text{Inf}$ , a domain error shall occur, and either a NaN (if supported), or an implementation-  
 46080 defined value shall be returned.

46081 If the correct value would cause underflow, and is representable, a range error may occur and  
 46082 the correct value shall be returned.

46083 XSI If the correct value would cause overflow, a range error shall occur and *tan()*, *tanf()*, and *tanl()*  
 46084 shall return the value of the macro HUGE\_VAL, HUGE\_VALF, and HUGE\_VALL, respectively.

## 46085 ERRORS

46086 These functions shall fail if:

46087 MX Domain Error The value *x* is  $\pm\text{Inf}$ .

46088 If the integer expression (math\_errhandling & MATH\_ERRNO) is non-zero, |  
 46089 then *errno* shall be set to [EDOM]. If the integer expression (math\_errhandling |  
 46090 & MATH\_ERREXCEPT) is non-zero, then the invalid floating-point exception |  
 46091 shall be raised. |

46092 XSI Range Error The result overflows

46093 If the integer expression (math\_errhandling & MATH\_ERRNO) is non-zero, |  
 46094 then *errno* shall be set to [ERANGE]. If the integer expression |  
 46095 (math\_errhandling & MATH\_ERREXCEPT) is non-zero, then the overflow |  
 46096 floating-point exception shall be raised. |

46097 These functions may fail if:

46098 MX      Range Error      The result underflows, or the value  $x$  is subnormal.

46099                              If the integer expression (`math_errhandling` & `MATH_ERRNO`) is non-zero, |  
 46100                              then *errno* shall be set to [ERANGE]. If the integer expression |  
 46101                              (`math_errhandling` & `MATH_ERREXCEPT`) is non-zero, then the underflow |  
 46102                              floating-point exception shall be raised. |

#### 46103 EXAMPLES

##### 46104              Taking the Tangent of a 45-Degree Angle

```
46105              #include <math.h>
46106              ...
46107              double radians = 45.0 * M_PI / 180;
46108              double result;
46109              ...
46110              result = tan (radians);
```

#### 46111 APPLICATION USAGE

46112              There are no known floating-point representations such that for a normal argument,  $\tan(x)$  is  
 46113              either overflow or underflow.

46114              These functions may lose accuracy when their argument is near a multiple of  $\pi/2$  or is far from  
 46115              0.0.

46116              On error, the expressions (`math_errhandling` & `MATH_ERRNO`) and (`math_errhandling` &  
 46117              `MATH_ERREXCEPT`) are independent of each other, but at least one of them must be non-zero.

#### 46118 RATIONALE

46119              None.

#### 46120 FUTURE DIRECTIONS

46121              None.

#### 46122 SEE ALSO

46123              *atan()*, *feclearexcept()*, *fetestexcept()*, *isnan()*, the Base Definitions volume of IEEE Std 1003.1-200x, |  
 46124              Section 4.18, Treatment of Error Conditions for Mathematical Functions, <**math.h**> |

#### 46125 CHANGE HISTORY

46126              First released in Issue 1. Derived from Issue 1 of the SVID.

#### 46127 Issue 5

46128              The last two paragraphs of the DESCRIPTION were included as APPLICATION USAGE notes  
 46129              in previous issues.

#### 46130 Issue 6

46131              The *tanf()* and *tanl()* functions are added for alignment with the ISO/IEC 9899:1999 standard.

46132              The DESCRIPTION, RETURN VALUE, ERRORS, and APPLICATION USAGE sections are  
 46133              revised to align with the ISO/IEC 9899:1999 standard.

46134              IEC 60559:1989 standard floating-point extensions over the ISO/IEC 9899:1999 standard are  
 46135              marked.

46136 **NAME**

46137           tanf — tangent function

46138 **SYNOPSIS**

46139           #include &lt;math.h&gt;

46140           float tanf(float x);

46141 **DESCRIPTION**46142           Refer to *tan()*.

46143 **NAME**

46144 tanh, tanhf, tanhl — hyperbolic tangent functions

46145 **SYNOPSIS**

46146 #include &lt;math.h&gt;

46147 double tanh(double x);

46148 float tanhf(float x);

46149 long double tanhl(long double x);

46150 **DESCRIPTION**

46151 **CX** The functionality described on this reference page is aligned with the ISO C standard. Any  
 46152 conflict between the requirements described here and the ISO C standard is unintentional. This  
 46153 volume of IEEE Std 1003.1-200x defers to the ISO C standard.

46154 These functions shall compute the hyperbolic tangent of their argument *x*.

46155 An application wishing to check for error situations should set *errno* to zero and call  
 46156 *feclearexcept*(FE\_ALL\_EXCEPT) before calling these functions. On return, if *errno* is non-zero or  
 46157 *fetestexcept*(FE\_INVALID | FE\_DIVBYZERO | FE\_OVERFLOW | FE\_UNDERFLOW) is non-  
 46158 zero, an error has occurred.

46159 **RETURN VALUE**46160 Upon successful completion, these functions shall return the hyperbolic tangent of *x*.46161 **MX** If *x* is NaN, a NaN shall be returned.46162 If *x* is  $\pm 0$ , *x* shall be returned.46163 If *x* is  $\pm\text{Inf}$ ,  $\pm 1$  shall be returned.46164 If *x* is subnormal, a range error may occur and *x* should be returned.46165 **ERRORS**

46166 These functions may fail if:

46167 **MX** **Range Error** The value of *x* is subnormal.

46168 If the integer expression (math\_errhandling & MATH\_ERRNO) is non-zero, |  
 46169 then *errno* shall be set to [ERANGE]. If the integer expression |  
 46170 (math\_errhandling & MATH\_ERREXCEPT) is non-zero, then the underflow |  
 46171 floating-point exception shall be raised. |

46172 **EXAMPLES**

46173 None.

46174 **APPLICATION USAGE**

46175 On error, the expressions (math\_errhandling & MATH\_ERRNO) and (math\_errhandling &  
 46176 MATH\_ERREXCEPT) are independent of each other, but at least one of them must be non-zero.

46177 **RATIONALE**

46178 None.

46179 **FUTURE DIRECTIONS**

46180 None.

46181 **SEE ALSO**

46182 *atanh()*, *feclearexcept()*, *fetestexcept()*, *isnan()*, *tan()*, the Base Definitions volume of |  
 46183 IEEE Std 1003.1-200x, Section 4.18, Treatment of Error Conditions for Mathematical Functions, |  
 46184 <math.h>

46185 **CHANGE HISTORY**

46186 First released in Issue 1. Derived from Issue 1 of the SVID.

46187 **Issue 5**

46188 The DESCRIPTION is updated to indicate how an application should check for an error. This  
46189 text was previously published in the APPLICATION USAGE section.

46190 **Issue 6**

46191 The *tanhf()* and *tanhl()* functions are added for alignment with the ISO/IEC 9899:1999 standard.

46192 The DESCRIPTION, RETURN VALUE, ERRORS, and APPLICATION USAGE sections are  
46193 revised to align with the ISO/IEC 9899:1999 standard.

46194 IEC 60559:1989 standard floating-point extensions over the ISO/IEC 9899:1999 standard are  
46195 marked.



46196 **NAME**

46197           tanl — tangent function

46198 **SYNOPSIS**

46199           #include <math.h>

46200           long double tanl(long double x);

46201 **DESCRIPTION**

46202           Refer to *tan()*.

46203 **NAME**

46204 tcdrain — wait for transmission of output

46205 **SYNOPSIS**

46206 #include &lt;termios.h&gt;

46207 int tcdrain(int *fildev*);46208 **DESCRIPTION**46209 The *tcdrain()* function shall block until all output written to the object referred to by *fildev* is |  
46210 transmitted. The *fildev* argument is an open file descriptor associated with a terminal.46211 Any attempts to use *tcdrain()* from a process which is a member of a background process group |  
46212 on a *fildev* associated with its controlling terminal, shall cause the process group to be sent a |  
46213 SIGTTOU signal. If the calling process is blocking or ignoring SIGTTOU signals, the process |  
46214 shall be allowed to perform the operation, and no signal is sent.46215 **RETURN VALUE**46216 Upon successful completion, 0 shall be returned. Otherwise, -1 shall be returned and *errno* set to |  
46217 indicate the error.46218 **ERRORS**46219 The *tcdrain()* function shall fail if:46220 [EBADF] The *fildev* argument is not a valid file descriptor.46221 [EINTR] A signal interrupted *tcdrain()*.46222 [ENOTTY] The file associated with *fildev* is not a terminal.46223 The *tcdrain()* function may fail if:46224 [EIO] The process group of the writing process is orphaned, and the writing process |  
46225 is not ignoring or blocking SIGTTOU.46226 **EXAMPLES**

46227 None.

46228 **APPLICATION USAGE**

46229 None.

46230 **RATIONALE**

46231 None.

46232 **FUTURE DIRECTIONS**

46233 None.

46234 **SEE ALSO**46235 *tflush()*, the Base Definitions volume of IEEE Std 1003.1-200x, <termios.h>, <unistd.h>, the Base |  
46236 Definitions volume of IEEE Std 1003.1-200x, Chapter 11, General Terminal Interface46237 **CHANGE HISTORY**

46238 First released in Issue 3.

46239 Entry included for alignment with the POSIX.1-1988 standard.

46240 **Issue 6**46241 The following new requirements on POSIX implementations derive from alignment with the |  
46242 Single UNIX Specification:

- 46243 • In the DESCRIPTION, the final paragraph is no longer conditional on |
- 
- 46244 \_POSIX\_JOB\_CONTROL. This is a FIPS requirement.

46245

- The [EIO] error is added.

46246 **NAME**

46247 tcflow — suspend or restart the transmission or reception of data

46248 **SYNOPSIS**

46249 #include &lt;termios.h&gt;

46250 int tcflow(int *fildev*, int *action*);46251 **DESCRIPTION**46252 The *tcflow()* function shall suspend or restart transmission or reception of data on the object  
46253 referred to by *fildev*, depending on the value of *action*. The *fildev* argument is an open file  
46254 descriptor associated with a terminal.

- 46255
- If *action* is TCOFF, output shall be suspended.
  - If *action* is TCOON, suspended output shall be restarted.
  - If *action* is TCIOFF, the system shall transmit a STOP character, which is intended to cause  
46258 the terminal device to stop transmitting data to the system.
  - If *action* is TCION, the system shall transmit a START character, which is intended to cause  
46260 the terminal device to start transmitting data to the system.

46261 The default on the opening of a terminal file is that neither its input nor its output are  
46262 suspended.46263 Attempts to use *tcflow()* from a process which is a member of a background process group on a  
46264 *fildev* associated with its controlling terminal, shall cause the process group to be sent a  
46265 SIGTTOU signal. If the calling process is blocking or ignoring SIGTTOU signals, the process  
46266 shall be allowed to perform the operation, and no signal is sent.46267 **RETURN VALUE**46268 Upon successful completion, 0 shall be returned. Otherwise, -1 shall be returned and *errno* set to  
46269 indicate the error.46270 **ERRORS**46271 The *tcflow()* function shall fail if:

- 46272 [EBADF] The
- fildev*
- argument is not a valid file descriptor.
- 
- 46273 [EINVAL] The
- action*
- argument is not a supported value.
- 
- 46274 [ENOTTY] The file associated with
- fildev*
- is not a terminal.

46275 The *tcflow()* function may fail if:

- 46276 [EIO] The process group of the writing process is orphaned, and the writing process
- 
- 46277 is not ignoring or blocking SIGTTOU.

46278 **EXAMPLES**

46279 None.

46280 **APPLICATION USAGE**

46281 None.

46282 **RATIONALE**

46283 None.

46284 **FUTURE DIRECTIONS**

46285 None.

46286 **SEE ALSO**

46287 *tcsendbreak()*, the Base Definitions volume of IEEE Std 1003.1-200x, <termios.h>, <unistd.h>, the  
46288 Base Definitions volume of IEEE Std 1003.1-200x, Chapter 11, General Terminal Interface

46289 **CHANGE HISTORY**

46290 First released in Issue 3.

46291 Entry included for alignment with the POSIX.1-1988 standard.

46292 **Issue 6**

46293 The following new requirements on POSIX implementations derive from alignment with the  
46294 Single UNIX Specification:

- 46295 • The [EIO] error is added.

46296 **NAME**

46297 tcflush — flush non-transmitted output data, non-read input data, or both

46298 **SYNOPSIS**

46299 #include &lt;termios.h&gt;

46300 int tcflush(int *fildev*, int *queue\_selector*);46301 **DESCRIPTION**46302 Upon successful completion, *tcflush()* shall discard data written to the object referred to by *fildev*  
46303 (an open file descriptor associated with a terminal) but not transmitted, or data received but not  
46304 read, depending on the value of *queue\_selector*:

- 46305
- If *queue\_selector* is TCIFLUSH, it shall flush data received but not read.
  - If *queue\_selector* is TCOFLUSH, it shall flush data written but not transmitted.
  - If *queue\_selector* is TCIOFLUSH, it shall flush both data received but not read and data  
46308 written but not transmitted.

46309 Attempts to use *tcflush()* from a process which is a member of a background process group on a  
46310 *fildev* associated with its controlling terminal shall cause the process group to be sent a SIGTTOU |  
46311 signal. If the calling process is blocking or ignoring SIGTTOU signals, the process shall be |  
46312 allowed to perform the operation, and no signal is sent. |46313 **RETURN VALUE**46314 Upon successful completion, 0 shall be returned. Otherwise, -1 shall be returned and *errno* set to  
46315 indicate the error.46316 **ERRORS**46317 The *tcflush()* function shall fail if:

- 46318 [EBADF] The
- fildev*
- argument is not a valid file descriptor.
- 
- 46319 [EINVAL] The
- queue\_selector*
- argument is not a supported value.
- 
- 46320 [ENOTTY] The file associated with
- fildev*
- is not a terminal.

46321 The *tcflush()* function may fail if:

- 46322 [EIO] The process group of the writing process is orphaned, and the writing process
- 
- 46323 is not ignoring or blocking SIGTTOU.

46324 **EXAMPLES**

46325 None.

46326 **APPLICATION USAGE**

46327 None.

46328 **RATIONALE**

46329 None.

46330 **FUTURE DIRECTIONS**

46331 None.

46332 **SEE ALSO**46333 *tcdrain()*, the Base Definitions volume of IEEE Std 1003.1-200x, <termios.h>, <unistd.h>, the  
46334 Base Definitions volume of IEEE Std 1003.1-200x, Chapter 11, General Terminal Interface

46335 **CHANGE HISTORY**

46336 First released in Issue 3.

46337 Entry included for alignment with the POSIX.1-1988 standard.

46338 **Issue 6**46339 The Open Group Corrigendum U035/1 is applied. In the ERRORS and APPLICATION USAGE  
46340 sections, references to *tcfow*() are replaced with *tcflush*().46341 The following new requirements on POSIX implementations derive from alignment with the  
46342 Single UNIX Specification:46343 • In the DESCRIPTION, the final paragraph is no longer conditional on  
46344 `_POSIX_JOB_CONTROL`. This is a FIPS requirement.

46345 • The [EIO] error is added.

## 46346 NAME

46347 tcgetattr — get the parameters associated with the terminal

## 46348 SYNOPSIS

46349 #include <termios.h>

46350 int tcgetattr(int *fildev*, struct termios \**termios\_p*);

## 46351 DESCRIPTION

46352 The *tcgetattr()* function shall get the parameters associated with the terminal referred to by *fildev*  
46353 and store them in the **termios** structure referenced by *termios\_p*. The *fildev* argument is an open  
46354 file descriptor associated with a terminal.

46355 The *termios\_p* argument is a pointer to a **termios** structure.

46356 The *tcgetattr()* operation is allowed from any process.

46357 If the terminal device supports different input and output baud rates, the baud rates stored in  
46358 the **termios** structure returned by *tcgetattr()* shall reflect the actual baud rates, even if they are  
46359 equal. If differing baud rates are not supported, the rate returned as the output baud rate shall be  
46360 the actual baud rate. If the terminal device does not support split baud rates, the input baud rate  
46361 stored in the **termios** structure shall be the output rate (as one of the symbolic values).

## 46362 RETURN VALUE

46363 Upon successful completion, 0 shall be returned. Otherwise, -1 shall be returned and *errno* set to  
46364 indicate the error.

## 46365 ERRORS

46366 The *tcgetattr()* function shall fail if:

46367 [EBADF] The *fildev* argument is not a valid file descriptor.

46368 [ENOTTY] The file associated with *fildev* is not a terminal.

## 46369 EXAMPLES

46370 None.

## 46371 APPLICATION USAGE

46372 None.

## 46373 RATIONALE

46374 Care must be taken when changing the terminal attributes. Applications should always do a  
46375 *tcgetattr()*, save the **termios** structure values returned, and then do a *tcsetattr()* changing only  
46376 the necessary fields. The application should use the values saved from the *tcgetattr()* to reset the  
46377 terminal state whenever it is done with the terminal. This is necessary because terminal  
46378 attributes apply to the underlying port and not to each individual open instance; that is, all  
46379 processes that have used the terminal see the latest attribute changes.

46380 A program that uses these functions should be written to catch all signals and take other  
46381 appropriate actions to ensure that when the program terminates, whether planned or not, the  
46382 terminal device's state is restored to its original state.

46383 Existing practice dealing with error returns when only part of a request can be honored is based  
46384 on calls to the *ioctl()* function. In historical BSD and System V implementations, the  
46385 corresponding *ioctl()* returns zero if the requested actions were semantically correct, even if  
46386 some of the requested changes could not be made. Many existing applications assume this  
46387 behavior and would no longer work correctly if the return value were changed from zero to -1  
46388 in this case.



46389 Note that either specification has a problem. When zero is returned, it implies everything  
46390 succeeded even if some of the changes were not made. When -1 is returned, it implies  
46391 everything failed even though some of the changes were made.

46392 Applications that need all of the requested changes made to work properly should follow  
46393 *tcsetattr()* with a call to *tcgetattr()* and compare the appropriate field values.

46394 **FUTURE DIRECTIONS**

46395 None.

46396 **SEE ALSO**

46397 *tcsetattr()*, the Base Definitions volume of IEEE Std 1003.1-200x, <**termios.h**>, the Base  
46398 Definitions volume of IEEE Std 1003.1-200x, Chapter 11, General Terminal Interface

46399 **CHANGE HISTORY**

46400 First released in Issue 3.

46401 Entry included for alignment with the POSIX.1-1988 standard.

46402 **Issue 6**

46403 In the DESCRIPTION, the rate returned as the input baud rate shall be the output rate.

46404 Previously, the number zero was also allowed but was obsolescent.

46405 **NAME**

46406 tcgetpgrp — get the foreground process group ID

46407 **SYNOPSIS**

46408 #include &lt;unistd.h&gt;

46409 pid\_t tcgetpgrp(int *fdes*);46410 **DESCRIPTION**46411 The *tcgetpgrp()* function shall return the value of the process group ID of the foreground process  
46412 group associated with the terminal.46413 If there is no foreground process group, *tcgetpgrp()* shall return a value greater than 1 that does  
46414 not match the process group ID of any existing process group.46415 The *tcgetpgrp()* function is allowed from a process that is a member of a background process  
46416 group; however, the information may be subsequently changed by a process that is a member of  
46417 a foreground process group.46418 **RETURN VALUE**46419 Upon successful completion, *tcgetpgrp()* shall return the value of the process group ID of the  
46420 foreground process associated with the terminal. Otherwise, -1 shall be returned and *errno* set to  
46421 indicate the error.46422 **ERRORS**46423 The *tcgetpgrp()* function shall fail if:46424 [EBADF] The *fdes* argument is not a valid file descriptor.46425 [ENOTTY] The calling process does not have a controlling terminal, or the file is not the  
46426 controlling terminal.46427 **EXAMPLES**

46428 None.

46429 **APPLICATION USAGE**

46430 None.

46431 **RATIONALE**

46432 None.

46433 **FUTURE DIRECTIONS**

46434 None.

46435 **SEE ALSO**46436 *setsid()*, *setpgid()*, *tcsetpgrp()*, the Base Definitions volume of IEEE Std 1003.1-200x,  
46437 <sys/types.h>, <unistd.h>46438 **CHANGE HISTORY**

46439 First released in Issue 3.

46440 Entry included for alignment with the POSIX.1-1988 standard.

46441 **Issue 6**

46442 In the SYNOPSIS, the optional include of the &lt;sys/types.h&gt; header is removed.

46443 The following new requirements on POSIX implementations derive from alignment with the  
46444 Single UNIX Specification:

- 46445 • The requirement to include <sys/types.h> has been removed. Although <sys/types.h> was
- 
- 46446 required for conforming implementations of previous POSIX specifications, it was not
- 
- 46447 required for UNIX applications.

46448  
46449

- In the DESCRIPTION, text previously conditional on support for `_POSIX_JOB_CONTROL` is now mandatory. This is a FIPS requirement.

46450 **NAME**

46451 tcgetsid — get process group ID for session leader for controlling terminal

46452 **SYNOPSIS**46453 XSI `#include <termios.h>`46454 `pid_t tcgetsid(int fildes);`

46455

46456 **DESCRIPTION**46457 The *tcgetsid()* function shall obtain the process group ID of the session for which the terminal  
46458 specified by *fildes* is the controlling terminal.46459 **RETURN VALUE**46460 Upon successful completion, *tcgetsid()* shall return the process group ID associated with the  
46461 terminal. Otherwise, a value of **(pid\_t)-1** shall be returned and *errno* set to indicate the error.46462 **ERRORS**46463 The *tcgetsid()* function shall fail if:46464 [EBADF] The *fildes* argument is not a valid file descriptor.46465 [ENOTTY] The calling process does not have a controlling terminal, or the file is not the  
46466 controlling terminal.46467 **EXAMPLES**

46468 None.

46469 **APPLICATION USAGE**

46470 None.

46471 **RATIONALE**

46472 None.

46473 **FUTURE DIRECTIONS**

46474 None.

46475 **SEE ALSO**46476 The Base Definitions volume of IEEE Std 1003.1-200x, **<termios.h>**46477 **CHANGE HISTORY**

46478 First released in Issue 4, Version 2.

46479 **Issue 5**

46480 Moved from X/OPEN UNIX extension to BASE.

46481 The [EACCES] error has been removed from the list of mandatory errors, and the description of  
46482 [ENOTTY] has been reworded.

46483 **NAME**

46484 tcsendbreak — send a “break” for a specific duration

46485 **SYNOPSIS**

46486 #include &lt;termios.h&gt;

46487 int tcsendbreak(int *fildev*, int *duration*);46488 **DESCRIPTION**

46489 If the terminal is using asynchronous serial data transmission, *tcsendbreak()* shall cause |  
 46490 transmission of a continuous stream of zero-valued bits for a specific duration. If *duration* is 0, it |  
 46491 shall cause transmission of zero-valued bits for at least 0,25 seconds, and not more than 0,5 |  
 46492 seconds. If *duration* is not 0, it shall send zero-valued bits for an implementation-defined period |  
 46493 of time.

46494 The *fildev* argument is an open file descriptor associated with a terminal. |

46495 If the terminal is not using asynchronous serial data transmission, it is implementation-defined |  
 46496 whether *tcsendbreak()* sends data to generate a break condition or returns without taking any |  
 46497 action.

46498 Attempts to use *tcsendbreak()* from a process which is a member of a background process group |  
 46499 on a *fildev* associated with its controlling terminal shall cause the process group to be sent a |  
 46500 SIGTTOU signal. If the calling process is blocking or ignoring SIGTTOU signals, the process |  
 46501 shall be allowed to perform the operation, and no signal is sent. |

46502 **RETURN VALUE**

46503 Upon successful completion, 0 shall be returned. Otherwise, -1 shall be returned and *errno* set to |  
 46504 indicate the error.

46505 **ERRORS**46506 The *tcsendbreak()* function shall fail if:46507 [EBADF] The *fildev* argument is not a valid file descriptor.46508 [ENOTTY] The file associated with *fildev* is not a terminal.46509 The *tcsendbreak()* function may fail if:

46510 [EIO] The process group of the writing process is orphaned, and the writing process |  
 46511 is not ignoring or blocking SIGTTOU.

46512 **EXAMPLES**

46513 None.

46514 **APPLICATION USAGE**

46515 None.

46516 **RATIONALE**

46517 None.

46518 **FUTURE DIRECTIONS**

46519 None.

46520 **SEE ALSO**

46521 The Base Definitions volume of IEEE Std 1003.1-200x, <termios.h>, <unistd.h>, the Base |  
 46522 Definitions volume of IEEE Std 1003.1-200x, Chapter 11, General Terminal Interface

46523 **CHANGE HISTORY**

46524 First released in Issue 3.

46525 Entry included for alignment with the POSIX.1-1988 standard.

46526 **Issue 6**46527 The following new requirements on POSIX implementations derive from alignment with the  
46528 Single UNIX Specification:

- 46529 • In the DESCRIPTION, text previously conditional on `_POSIX_JOB_CONTROL` is now
- 46530 mandated. This is a FIPS requirement.
- 46531 • The [EIO] error is added.

46532 **NAME**

46533 tcsetattr — set the parameters associated with the terminal

46534 **SYNOPSIS**

46535 #include &lt;termios.h&gt;

46536 int tcsetattr(int *fildev*, int *optional\_actions*,  
46537 const struct termios \**termios\_p*);46538 **DESCRIPTION**46539 The *tcsetattr()* function shall set the parameters associated with the terminal referred to by the  
46540 open file descriptor *fildev* (an open file descriptor associated with a terminal) from the **termios**  
46541 structure referenced by *termios\_p* as follows:

- 46542
- If *optional\_actions* is TCSANOW, the change shall occur immediately.
  - If *optional\_actions* is TCSADRAIN, the change shall occur after all output written to *fildev* is  
46543 transmitted. This function should be used when changing parameters that affect output.
  - If *optional\_actions* is TCSAFLUSH, the change shall occur after all output written to *fildev* is  
46544 transmitted, and all input so far received but not read shall be discarded before the change is  
46545 made.

46546 If the output baud rate stored in the **termios** structure pointed to by *termios\_p* is the zero baud  
46547 rate, B0, the modem control lines shall no longer be asserted. Normally, this shall disconnect the  
46548 line.46549 If the input baud rate stored in the **termios** structure pointed to by *termios\_p* is 0, the input baud  
46550 rate given to the hardware is the same as the output baud rate stored in the **termios** structure.46551 The *tcsetattr()* function shall return successfully if it was able to perform any of the requested  
46552 actions, even if some of the requested actions could not be performed. It shall set all the  
46553 attributes that the implementation supports as requested and leaves all the attributes not  
46554 supported by the implementation unchanged. If no part of the request can be honored, it shall  
46555 return  $-1$  and set *errno* to [EINVAL]. If the input and output baud rates differ and are a  
46556 combination that is not supported, neither baud rate shall be changed. A subsequent call to  
46557 *tcgetattr()* shall return the actual state of the terminal device (reflecting both the changes made  
46558 and not made in the previous *tcsetattr()* call). The *tcsetattr()* function shall not change the values  
46559 found in the **termios** structure under any circumstances.46560 The effect of *tcsetattr()* is undefined if the value of the **termios** structure pointed to by *termios\_p*  
46561 was not derived from the result of a call to *tcgetattr()* on *fildev*; an application should modify  
46562 only fields and flags defined by this volume of IEEE Std 1003.1-200x between the call to  
46563 *tcgetattr()* and *tcsetattr()*, leaving all other fields and flags unmodified.46564 No actions defined by this volume of IEEE Std 1003.1-200x, other than a call to *tcsetattr()* or a  
46565 close of the last file descriptor in the system associated with this terminal device, shall cause any  
46566 of the terminal attributes defined by this volume of IEEE Std 1003.1-200x to change.46567 If *tcsetattr()* is called from a process which is a member of a background process group on a  
46568 *fildev* associated with its controlling terminal:

- 46569
- If the calling process is blocking or ignoring SIGTTOU signals, the operation completes  
46570 normally and no signal is sent.
  - Otherwise, a SIGTTOU signal shall be sent to the process group.

46574 **RETURN VALUE**

46575 Upon successful completion, 0 shall be returned. Otherwise, -1 shall be returned and  
46576 *errno* set to indicate the error.

46577 **ERRORS**

46578 The *tcsetattr()* function shall fail if:

46579 [EBADF] The *fildev* argument is not a valid file descriptor.

46580 [EINTR] A signal interrupted *tcsetattr()*.

46581 [EINVAL] The *optional\_actions* argument is not a supported value, or an attempt was  
46582 made to change an attribute represented in the **termios** structure to an  
46583 unsupported value.

46584 [ENOTTY] The file associated with *fildev* is not a terminal.

46585 The *tcsetattr()* function may fail if:

46586 [EIO] The process group of the writing process is orphaned, and the writing process  
46587 is not ignoring or blocking SIGTTOU.

46588 **EXAMPLES**

46589 None.

46590 **APPLICATION USAGE**

46591 If trying to change baud rates, applications should call *tcsetattr()* then call *tcgetattr()* in order to  
46592 determine what baud rates were actually selected.

46593 **RATIONALE**

46594 The *tcsetattr()* function can be interrupted in the following situations:

- 46595 • It is interrupted while waiting for output to drain.
- 46596 • It is called from a process in a background process group and SIGTTOU is caught.

46597 See also the RATIONALE section in *tcgetattr()*.

46598 **FUTURE DIRECTIONS**

46599 Using an input baud rate of 0 to set the input rate equal to the output rate may not necessarily be  
46600 supported in a future version of this volume of IEEE Std 1003.1-200x.

46601 **SEE ALSO**

46602 *cfgetispeed()*, *tcgetattr()*, the Base Definitions volume of IEEE Std 1003.1-200x, <**termios.h**>,  
46603 <**unistd.h**>, the Base Definitions volume of IEEE Std 1003.1-200x, Chapter 11, General Terminal  
46604 Interface

46605 **CHANGE HISTORY**

46606 First released in Issue 3.

46607 Entry included for alignment with the POSIX.1-1988 standard.

46608 **Issue 6**

46609 The following new requirements on POSIX implementations derive from alignment with the  
46610 Single UNIX Specification:

- 46611 • In the DESCRIPTION, text previously conditional on `_POSIX_JOB_CONTROL` is now  
46612 mandated. This is a FIPS requirement.
- 46613 • The [EIO] error is added.



46614  
46615

In the DESCRIPTION, the text describing use of *tcsetattr()* from a process which is a member of a background process group is clarified.

46616 **NAME**

46617 tcsetpgrp — set the foreground process group ID

46618 **SYNOPSIS**

46619 #include &lt;unistd.h&gt;

46620 int tcsetpgrp(int *fildes*, pid\_t *pgid\_id*);46621 **DESCRIPTION**

46622 If the process has a controlling terminal, *tcsetpgrp()* shall set the foreground process group ID  
46623 associated with the terminal to *pgid\_id*. The application shall ensure that the file associated with  
46624 *fildes* is the controlling terminal of the calling process and the controlling terminal is currently  
46625 associated with the session of the calling process. The application shall ensure that the value of  
46626 *pgid\_id* matches a process group ID of a process in the same session as the calling process.

46627 Attempts to use *tcsetpgrp()* from a process which is a member of a background process group on  
46628 a *fildes* associated with its controlling terminal shall cause the process group to be sent a  
46629 SIGTTOU signal. If the calling process is blocking or ignoring SIGTTOU signals, the process  
46630 shall be allowed to perform the operation, and no signal is sent.

46631 **RETURN VALUE**

46632 Upon successful completion, 0 shall be returned. Otherwise, -1 shall be returned and *errno* set to  
46633 indicate the error.

46634 **ERRORS**46635 The *tcsetpgrp()* function shall fail if:

46636 [EBADF] The *fildes* argument is not a valid file descriptor.

46637 [EINVAL] This implementation does not support the value in the *pgid\_id* argument.

46638 [ENOTTY] The calling process does not have a controlling terminal, or the file is not the  
46639 controlling terminal, or the controlling terminal is no longer associated with  
46640 the session of the calling process.

46641 [EPERM] The value of *pgid\_id* is a value supported by the implementation, but does not  
46642 match the process group ID of a process in the same session as the calling  
46643 process.

46644 **EXAMPLES**

46645 None.

46646 **APPLICATION USAGE**

46647 None.

46648 **RATIONALE**

46649 None.

46650 **FUTURE DIRECTIONS**

46651 None.

46652 **SEE ALSO**46653 *tcgetpgrp()*, the Base Definitions volume of IEEE Std 1003.1-200x, <sys/types.h>, <unistd.h>46654 **CHANGE HISTORY**

46655 First released in Issue 3.

46656 Entry included for alignment with the POSIX.1-1988 standard.

46657 **Issue 6**

46658 In the SYNOPSIS, the inclusion of `<sys/types.h>` is no longer required.

46659 The following new requirements on POSIX implementations derive from alignment with the  
46660 Single UNIX Specification:

46661 • The requirement to include `<sys/types.h>` has been removed. Although `<sys/types.h>` was  
46662 required for conforming implementations of previous POSIX specifications, it was not  
46663 required for UNIX applications.

46664 • In the DESCRIPTION and ERRORS sections, text previously conditional on  
46665 `_POSIX_JOB_CONTROL` is now mandated. This is a FIPS requirement.

46666 The DESCRIPTION is updated to avoid use of the term “must” for application requirements.

46667 The Open Group Corrigendum U047/4 is applied.

## 46668 NAME

46669 tdelete, tfind, tsearch, twalk — manage a binary search tree

## 46670 SYNOPSIS

```

46671 xsi #include <search.h>
46672 void *tdelete(const void *restrict key, void **restrict rootp,
46673             int(*compar)(const void *, const void *));
46674 void *tfind(const void *key, void *const *rootp,
46675            int(*compar)(const void *, const void *));
46676 void *tsearch(const void *key, void **rootp,
46677              int (*compar)(const void *, const void *));
46678 void twalk(const void *root,
46679           void (*action)(const void *, VISIT, int));
46680

```

## 46681 DESCRIPTION

46682 The *tdelete()*, *tfind()*, *tsearch()*, and *twalk()* functions manipulate binary search trees.  
 46683 Comparisons are made with a user-supplied routine, the address of which is passed as the  
 46684 *compar* argument. This routine is called with two arguments, the pointers to the elements being  
 46685 compared. The application shall ensure that the user-supplied routine returns an integer less  
 46686 than, equal to, or greater than 0, according to whether the first argument is to be considered less  
 46687 than, equal to, or greater than the second argument. The comparison function need not compare  
 46688 every byte, so arbitrary data may be contained in the elements in addition to the values being  
 46689 compared.

46690 The *tsearch()* function shall build and access the tree. The *key* argument is a pointer to an element |  
 46691 to be accessed or stored. If there is a node in the tree whose element is equal to the value pointed |  
 46692 to by *key*, a pointer to this found node shall be returned. Otherwise, the value pointed to by *key* |  
 46693 shall be inserted (that is, a new node is created and the value of *key* is copied to this node), and a |  
 46694 pointer to this node returned. Only pointers are copied, so the application shall ensure that the |  
 46695 calling routine stores the data. The *rootp* argument points to a variable that points to the root |  
 46696 node of the tree. A null pointer value for the variable pointed to by *rootp* denotes an empty tree; |  
 46697 in this case, the variable shall be set to point to the node which shall be at the root of the new |  
 46698 tree.

46699 Like *tsearch()*, *tfind()* shall search for a node in the tree, returning a pointer to it if found.  
 46700 However, if it is not found, *tfind()* shall return a null pointer. The arguments for *tfind()* are the  
 46701 same as for *tsearch()*.

46702 The *tdelete()* function shall delete a node from a binary search tree. The arguments are the same |  
 46703 as for *tsearch()*. The variable pointed to by *rootp* shall be changed if the deleted node was the |  
 46704 root of the tree. The *tdelete()* function shall return a pointer to the parent of the deleted node, or a |  
 46705 null pointer if the node is not found.

46706 The *twalk()* function shall traverse a binary search tree. The *root* argument is a pointer to the root |  
 46707 node of the tree to be traversed. (Any node in a tree may be used as the root for a walk below |  
 46708 that node.) The argument *action* is the name of a routine to be invoked at each node. This routine |  
 46709 is, in turn, called with three arguments. The first argument shall be the address of the node being |  
 46710 visited. The structure pointed to by this argument is unspecified and shall not be modified by |  
 46711 the application, but it shall be possible to cast a pointer-to-node into a pointer-to-pointer-to- |  
 46712 element to access the element stored in the node. The second argument shall be a value from an |  
 46713 enumeration data type:

```

46714 typedef enum { preorder, postorder, endorder, leaf } VISIT;

```

46715 (defined in `<search.h>`), depending on whether this is the first, second, or third time that the  
 46716 node is visited (during a depth-first, left-to-right traversal of the tree), or whether the node is a  
 46717 leaf. The third argument shall be the level of the node in the tree, with the root being level 0.

46718 If the calling function alters the pointer to the root, the result is undefined.

#### 46719 RETURN VALUE

46720 If the node is found, both `tsearch()` and `tfind()` shall return a pointer to it. If not, `tfind()` shall  
 46721 return a null pointer, and `tsearch()` shall return a pointer to the inserted item.

46722 A null pointer shall be returned by `tsearch()` if there is not enough space available to create a new  
 46723 node.

46724 A null pointer shall be returned by `tdelete()`, `tfind()`, and `tsearch()` if `rootp` is a null pointer on  
 46725 entry.

46726 The `tdelete()` function shall return a pointer to the parent of the deleted node, or a null pointer if  
 46727 the node is not found.

46728 The `twalk()` function shall not return a value.

#### 46729 ERRORS

46730 No errors are defined.

#### 46731 EXAMPLES

46732 The following code reads in strings and stores structures containing a pointer to each string and  
 46733 a count of its length. It then walks the tree, printing out the stored strings and their lengths in  
 46734 alphabetical order.

```

46735 #include <search.h>
46736 #include <string.h>
46737 #include <stdio.h>

46738 #define STRSZ    10000
46739 #define NODSZ    500

46740 struct node {          /* Pointers to these are stored in the tree. */
46741     char    *string;
46742     int     length;
46743 };

46744 char    string_space[STRSZ]; /* Space to store strings. */
46745 struct node nodes[NODSZ];   /* Nodes to store. */
46746 void    *root = NULL;      /* This points to the root. */

46747 int main(int argc, char *argv[])
46748 {
46749     char    *strptr = string_space;
46750     struct node *nodeptr = nodes;
46751     void    print_node(const void *, VISIT, int);
46752     int     i = 0, node_compare(const void *, const void *);

46753     while (gets(strptr) != NULL && i++ < NODSZ) {
46754         /* Set node. */
46755         nodeptr->string = strptr;
46756         nodeptr->length = strlen(strptr);
46757         /* Put node into the tree. */
46758         (void) tsearch((void *)nodeptr, (void **)&root,
46759             node_compare);

```

```

46760         /* Adjust pointers, so we do not overwrite tree. */
46761         strptr += nodeptr->length + 1;
46762         nodeptr++;
46763     }
46764     twalk(root, print_node);
46765     return 0;
46766 }
46767 /*
46768  * This routine compares two nodes, based on an
46769  * alphabetical ordering of the string field.
46770  */
46771 int
46772 node_compare(const void *node1, const void *node2)
46773 {
46774     return strcmp(((const struct node *) node1)->string,
46775                 ((const struct node *) node2)->string);
46776 }
46777 /*
46778  * This routine prints out a node, the second time
46779  * twalk encounters it or if it is a leaf.
46780  */
46781 void
46782 print_node(const void *ptr, VISIT order, int level)
46783 {
46784     const struct node *p = *(const struct node **) ptr;
46785     if (order == postorder || order == leaf) {
46786         (void) printf("string = %s, length = %d\n",
46787                     p->string, p->length);
46788     }
46789 }

```

#### 46790 APPLICATION USAGE

46791 The *root* argument to *twalk()* is one level of indirection less than the *rootp* arguments to *tdelete()*  
 46792 and *tsearch()*.

46793 There are two nomenclatures used to refer to the order in which tree nodes are visited. The  
 46794 *tsearch()* function uses **preorder**, **postorder**, and **endorder** to refer respectively to visiting a node  
 46795 before any of its children, after its left child and before its right, and after both its children. The  
 46796 alternative nomenclature uses **preorder**, **inorder**, and **postorder** to refer to the same visits, which  
 46797 could result in some confusion over the meaning of **postorder**.

#### 46798 RATIONALE

46799 None.

#### 46800 FUTURE DIRECTIONS

46801 None.

#### 46802 SEE ALSO

46803 *hcreate()*, *tsearch()*, the Base Definitions volume of IEEE Std 1003.1-200x, <[search.h](#)>

46804 **CHANGE HISTORY**

46805 First released in Issue 1. Derived from Issue 1 of the SVID.

46806 **Issue 5**

46807 The last paragraph of the DESCRIPTION was included as an APPLICATION USAGE note in  
46808 previous issues.

46809 **Issue 6**

46810 The DESCRIPTION is updated to avoid use of the term “must” for application requirements.

46811 The **restrict** keyword is added to the *tdelete()* prototype for alignment with the  
46812 ISO/IEC 9899:1999 standard.

46813 **NAME**

46814 tellmdir — current location of a named directory stream

46815 **SYNOPSIS**

46816 XSI #include &lt;dirent.h&gt;

46817 long tellmdir(DIR \*dirp);

46818

46819 **DESCRIPTION**46820 The *tellmdir()* function shall obtain the current location associated with the directory stream |  
46821 specified by *dirp*. |46822 If the most recent operation on the directory stream was a *seekdir()*, the directory position  
46823 returned from the *tellmdir()* shall be the same as that supplied as a *loc* argument for *seekdir()*.46824 **RETURN VALUE**46825 Upon successful completion, *tellmdir()* shall return the current location of the specified directory  
46826 stream.46827 **ERRORS**

46828 No errors are defined.

46829 **EXAMPLES**

46830 None.

46831 **APPLICATION USAGE**

46832 None.

46833 **RATIONALE**

46834 None.

46835 **FUTURE DIRECTIONS**

46836 None.

46837 **SEE ALSO**46838 *opendir()*, *readdir()*, *seekdir()*, the Base Definitions volume of IEEE Std 1003.1-200x, <dirent.h>46839 **CHANGE HISTORY**

46840 First released in Issue 2.



46841 **NAME**

46842 tempnam — create a name for a temporary file

46843 **SYNOPSIS**46844 XSI `#include <stdio.h>`46845 `char *tempnam(const char *dir, const char *pfx);`

46846

46847 **DESCRIPTION**46848 The *tempnam()* function shall generate a pathname that may be used for a temporary file. |

46849 The *tempnam()* function allows the user to control the choice of a directory. The *dir* argument  
46850 points to the name of the directory in which the file is to be created. If *dir* is a null pointer or  
46851 points to a string which is not a name for an appropriate directory, the path prefix defined as  
46852 P\_tmpdir in the <stdio.h> header shall be used. If that directory is not accessible, an  
46853 implementation-defined directory may be used.

46854 Many applications prefer their temporary files to have certain initial letter sequences in their  
46855 names. The *pfx* argument should be used for this. This argument may be a null pointer or point  
46856 to a string of up to five bytes to be used as the beginning of the filename.

46857 Some implementations of *tempnam()* may use *tmpnam()* internally. On such implementations, if  
46858 called more than {TMP\_MAX} times in a single process, the behavior is implementation-defined.

46859 **RETURN VALUE**

46860 Upon successful completion, *tempnam()* shall allocate space for a string, put the generated |  
46861 pathname in that space, and return a pointer to it. The pointer shall be suitable for use in a |  
46862 subsequent call to *free()*. Otherwise, it shall return a null pointer and set *errno* to indicate the  
46863 error.

46864 **ERRORS**46865 The *tempnam()* function shall fail if:

46866 [ENOMEM] Insufficient storage space is available.

46867 **EXAMPLES**46868 **Generating a Pathname** |

46869 The following example generates a pathname for a temporary file in directory **/tmp**, with the |  
46870 prefix *file*. After the filename has been created, the call to *free()* deallocates the space used to  
46871 store the filename.

```
46872 #include <stdio.h>
46873 #include <stdlib.h>
46874 ...
46875 char *directory = "/tmp";
46876 char *fileprefix = "file";
46877 char *file;

46878 file = tempnam(directory, fileprefix);
46879 free(file);
```

46880 **APPLICATION USAGE**

46881 This function only creates pathnames. It is the application's responsibility to create and remove |  
46882 the files. Between the time a pathname is created and the file is opened, it is possible for some |  
46883 other process to create a file with the same name. Applications may find *tmpfile()* more useful.

46884 **RATIONALE**

46885           None.

46886 **FUTURE DIRECTIONS**

46887           None.

46888 **SEE ALSO**

46889           *fopen()*, *free()*, *open()*, *tmpfile()*, *tmpnam()*, *unlink()*, the Base Definitions volume of  
46890           IEEE Std 1003.1-200x, <**stdio.h**>

46891 **CHANGE HISTORY**

46892           First released in Issue 1. Derived from Issue 1 of the SVID.

46893 **Issue 5**

46894           The last paragraph of the DESCRIPTION was included as an APPLICATION USAGE note in  
46895           previous issues.

46896 **NAME**

46897       tfind — search binary search tree

46898 **SYNOPSIS**

46899 XSI       #include &lt;search.h&gt;

46900       void \*tfind(const void \*key, void \*const \*rootp,  
46901                 int (\*compar)(const void \*, const void \*));

46902

46903 **DESCRIPTION**46904       Refer to *tdelete()*.

46905 **NAME**

46906 tgamma, tgammaf, tgammaL — compute gamma() function

46907 **SYNOPSIS**

```
46908 #include <math.h>
46909 double tgamma(double x);
46910 float tgammaf(float x);
46911 long double tgammaL(long double x);
```

46912 **DESCRIPTION**

46913 **CX** The functionality described on this reference page is aligned with the ISO C standard. Any  
 46914 conflict between the requirements described here and the ISO C standard is unintentional. This  
 46915 volume of IEEE Std 1003.1-200x defers to the ISO C standard.

46916 These functions shall compute the *gamma()* function of *x*.

46917 An application wishing to check for error situations should set *errno* to zero and call  
 46918 *feclearexcept*(FE\_ALL\_EXCEPT) before calling these functions. On return, if *errno* is non-zero or  
 46919 *fetestexcept*(FE\_INVALID | FE\_DIVBYZERO | FE\_OVERFLOW | FE\_UNDERFLOW) is non-  
 46920 zero, an error has occurred.

46921 **RETURN VALUE**

46922 Upon successful completion, these functions shall return *Gamma(x)*.

46923 If *x* is a negative integer, a domain error shall occur, and either a NaN (if supported), or an  
 46924 implementation-defined value shall be returned.

46925 If the correct value would cause overflow, a range error shall occur and *tgamma()*, *tgammaf()*,  
 46926 and *tgammaL()* shall return the value of the macro HUGE\_VAL, HUGE\_VALF, or HUGE\_VALL,  
 46927 respectively.

46928 **MX** If *x* is NaN, a NaN shall be returned.

46929 If *x* is +Inf, *x* shall be returned.

46930 If *x* is ±0, a pole error shall occur, and *tgamma()*, *tgammaf()*, and *tgammaL()* shall return  
 46931 ±HUGE\_VAL, ±HUGE\_VALF, and ±HUGE\_VALL, respectively.

46932 If *x* is -Inf, a domain error shall occur, and either a NaN (if supported), or an implementation-  
 46933 defined value shall be returned.

46934 **ERRORS**

46935 These functions shall fail if:

46936 **MX** Domain Error The value of *x* is a negative integer, or *x* is -Inf.

46937 If the integer expression (math\_errhandling & MATH\_ERRNO) is non-zero, |  
 46938 then *errno* shall be set to [EDOM]. If the integer expression (math\_errhandling |  
 46939 & MATH\_ERREXCEPT) is non-zero, then the invalid floating-point exception |  
 46940 shall be raised. |

46941 **MX** Pole Error The value of *x* is zero.

46942 If the integer expression (math\_errhandling & MATH\_ERRNO) is non-zero, |  
 46943 then *errno* shall be set to [ERANGE]. If the integer expression |  
 46944 (math\_errhandling & MATH\_ERREXCEPT) is non-zero, then the divide-by- |  
 46945 zero floating-point exception shall be raised. |

46946 Range Error The value overflows.

46947 If the integer expression (`math_errhandling & MATH_ERRNO`) is non-zero, |  
46948 then *errno* shall be set to [ERANGE]. If the integer expression |  
46949 (`math_errhandling & MATH_ERREXCEPT`) is non-zero, then the overflow |  
46950 floating-point exception shall be raised. |

**46951 EXAMPLES**

46952 None.

**46953 APPLICATION USAGE**

46954 For IEEE Std 754-1985 **double**, overflow happens when  $0 < x < 1/\text{DBL\_MAX}$ , and  $171.7 < x$ .  
46955 Overflow also happens near negative integers.

46956 On error, the expressions (`math_errhandling & MATH_ERRNO`) and (`math_errhandling &`  
46957 `MATH_ERREXCEPT`) are independent of each other, but at least one of them must be non-zero.

**46958 RATIONALE**

46959 This function is named *tgamma()* in order to avoid conflicts with the historical *gamma()* and  
46960 *lgamma()* functions.

**46961 FUTURE DIRECTIONS**

46962 It is possible that the error response for a negative integer argument may be changed to a pole  
46963 error and a return value of  $\pm\text{Inf}$ .

**46964 SEE ALSO**

46965 *feclearexcept()*, *fetestexcept()*, *lgamma()*, the Base Definitions volume of IEEE Std 1003.1-200x, |  
46966 Section 4.18, Treatment of Error Conditions for Mathematical Functions, <math.h> |

**46967 CHANGE HISTORY**

46968 First released in Issue 6. Derived from the ISO/IEC 9899:1999 standard.

46969 **NAME**

46970 time — get time

46971 **SYNOPSIS**

46972 #include &lt;time.h&gt;

46973 time\_t time(time\_t \*tloc);

46974 **DESCRIPTION**

46975 CX The functionality described on this reference page is aligned with the ISO C standard. Any  
46976 conflict between the requirements described here and the ISO C standard is unintentional. This  
46977 volume of IEEE Std 1003.1-200x defers to the ISO C standard.

46978 CX The *time()* function shall return the value of time in seconds since the Epoch.

46979 The *tloc* argument points to an area where the return value is also stored. If *tloc* is a null pointer,  
46980 no value is stored.

46981 **RETURN VALUE**

46982 Upon successful completion, *time()* shall return the value of time. Otherwise, (**time\_t**)−1 shall be  
46983 returned.

46984 **ERRORS**

46985 No errors are defined.

46986 **EXAMPLES**46987 **Getting the Current Time**

46988 The following example uses the *time()* function to calculate the time elapsed, in seconds, since  
46989 January 1, 1970 0:00 UTC, *localtime()* to convert that value to a broken-down time, and *asctime()*  
46990 to convert the broken-down time values into a printable string.

46991 #include &lt;stdio.h&gt;

46992 #include &lt;time.h&gt;

46993 main()

46994 {

46995 time\_t result;

46996 result = time(NULL);

46997 printf("%s%ld secs since the Epoch\n",

46998 asctime(localtime(&amp;result)),

46999 (long)result);

47000 return(0);

47001 }

47002 This example writes the current time to *stdout* in a form like this:

47003 Wed Jun 26 10:32:15 1996

47004 835810335 secs since the Epoch

47005 **Timing an Event**

47006 The following example gets the current time, prints it out in the user's format, and prints the  
47007 number of minutes to an event being timed.

```
47008 #include <time.h>
47009 #include <stdio.h>
47010 ...
47011 time_t now;
47012 int minutes_to_event;
47013 ...
47014 time(&now);
47015 minutes_to_event = ...;
47016 printf("The time is ");
47017 puts(asctime(localtime(&now)));
47018 printf("There are %d minutes to the event.\n",
47019         minutes_to_event);
47020 ...
```

47021 **APPLICATION USAGE**

47022 None.

47023 **RATIONALE**

47024 The *time()* function returns a value in seconds (type **time\_t**) while *times()* returns a set of values  
47025 in clock ticks (type **clock\_t**). Some historical implementations, such as 4.3 BSD, have  
47026 mechanisms capable of returning more precise times (see below). A generalized timing scheme  
47027 to unify these various timing mechanisms has been proposed but not adopted.

47028 Implementations in which **time\_t** is a 32-bit signed integer (many historical implementations)  
47029 fail in the year 2038. IEEE Std 1003.1-200x does not address this problem. However, the use of  
47030 the **time\_t** type is mandated in order to ease the eventual fix.

47031 The use of the **<time.h>**, header instead of **<sys/types.h>**, allows compatibility with the ISO C  
47032 standard.

47033 Many historical implementations (including Version 7) and the 1984 /usr/group standard use  
47034 **long** instead of **time\_t**. This volume of IEEE Std 1003.1-200x uses the latter type in order to agree  
47035 with the ISO C standard.

47036 4.3 BSD includes *time()* only as an alternate function to the more flexible *gettimeofday()* function.

47037 **FUTURE DIRECTIONS**

47038 In a future version of this volume of IEEE Std 1003.1-200x, **time\_t** is likely to be required to be  
47039 capable of representing times far in the future. Whether this will be mandated as a 64-bit type or  
47040 a requirement that a specific date in the future be representable (for example, 10000 AD) is not  
47041 yet determined. Systems purchased after the approval of this volume of IEEE Std 1003.1-200x  
47042 should be evaluated to determine whether their lifetime will extend past 2038.

47043 **SEE ALSO**

47044 *asctime()*, *clock()*, *ctime()*, *difftime()*, *gmtime()*, *localtime()*, *mktime()*, *strptime()*, *strptime()*, *utime()*,  
47045 the Base Definitions volume of IEEE Std 1003.1-200x, **<time.h>**

47046 **CHANGE HISTORY**

47047 First released in Issue 1. Derived from Issue 1 of the SVID.

47048 **Issue 6**

47049 Extensions beyond the ISO C standard are now marked. |

47050 The EXAMPLES, RATIONALE, and FUTURE DIRECTIONS sections are added. |



## 47051 NAME

47052 timer\_create — create a per-process timer (**REALTIME**)

## 47053 SYNOPSIS

47054 TMR #include &lt;signal.h&gt;

47055 #include &lt;time.h&gt;

47056 int timer\_create(clockid\_t *clockid*, struct sigevent \*restrict *evp*,47057 timer\_t \*restrict *timerid*);

47058

## 47059 DESCRIPTION

47060 The *timer\_create()* function shall create a per-process timer using the specified clock, *clock\_id*, as  
 47061 the timing base. The *timer\_create()* function shall return, in the location referenced by *timerid*, a  
 47062 timer ID of type **timer\_t** used to identify the timer in timer requests. This timer ID shall be  
 47063 unique within the calling process until the timer is deleted. The particular clock, *clock\_id*, is  
 47064 defined in <**time.h**>. The timer whose ID is returned shall be in a disarmed state upon return  
 47065 from *timer\_create()*.

47066 The *evp* argument, if non-NULL, points to a **sigevent** structure. This structure, allocated by the  
 47067 application, defines the asynchronous notification to occur as specified in Section 2.4.1 (on page  
 47068 478) when the timer expires. If the *evp* argument is NULL, the effect is as if the *evp* argument  
 47069 pointed to a **sigevent** structure with the *sigev\_notify* member having the value SIGEV\_SIGNAL,  
 47070 the *sigev\_signo* having a default signal number, and the *sigev\_value* member having the value of  
 47071 the timer ID.

47072 Each implementation shall define a set of clocks that can be used as timing bases for per-process  
 47073 MON timers. All implementations shall support a *clock\_id* of CLOCK\_REALTIME. If the Monotonic  
 47074 Clock option is supported, implementations shall support a *clock\_id* of CLOCK\_MONOTONIC.

47075 Per-process timers shall not be inherited by a child process across a *fork()* and shall be disarmed  
 47076 and deleted by an *exec*.

47077 CPT If \_POSIX\_CPUTIME is defined, implementations shall support *clock\_id* values representing the  
 47078 CPU-time clock of the calling process.

47079 TCT If \_POSIX\_THREAD\_CPUTIME is defined, implementations shall support *clock\_id* values  
 47080 representing the CPU-time clock of the calling thread.

47081 CPT|TCT It is implementation-defined whether a *timer\_create()* function will succeed if the value defined  
 47082 by *clock\_id* corresponds to the CPU-time clock of a process or thread different from the process  
 47083 or thread invoking the function.

## 47084 RETURN VALUE

47085 If the call succeeds, *timer\_create()* shall return zero and update the location referenced by *timerid*  
 47086 to a **timer\_t**, which can be passed to the per-process timer calls. If an error occurs, the function  
 47087 shall return a value of -1 and set *errno* to indicate the error. The value of *timerid* is undefined if  
 47088 an error occurs.

## 47089 ERRORS

47090 The *timer\_create()* function shall fail if:

47091 [EAGAIN] The system lacks sufficient signal queuing resources to honor the request.

47092 [EAGAIN] The calling process has already created all of the timers it is allowed by this  
47093 implementation.

47094 [EINVAL] The specified clock ID is not defined.

47095 CPT|TCT [ENOTSUP] The implementation does not support the creation of a timer attached to the  
47096 CPU-time clock that is specified by *clock\_id* and associated with a process or  
47097 thread different from the process or thread invoking *timer\_create()*.

47098 **EXAMPLES**

47099 None.

47100 **APPLICATION USAGE**

47101 None.

47102 **RATIONALE**

47103 **Periodic Timer Overrun and Resource Allocation**

47104 The specified timer facilities may deliver realtime signals (that is, queued signals) on |  
47105 implementations that support this option. Since realtime applications cannot afford to lose |  
47106 notifications of asynchronous events, like timer expirations or asynchronous I/O completions, it |  
47107 must be possible to ensure that sufficient resources exist to deliver the signal when the event |  
47108 occurs. In general, this is not a difficulty because there is a one-to-one correspondence between a |  
47109 request and a subsequent signal generation. If the request cannot allocate the signal delivery |  
47110 resources, it can fail the call with an [EAGAIN] error.

47111 Periodic timers are a special case. A single request can generate an unspecified number of |  
47112 signals. This is not a problem if the requesting process can service the signals as fast as they are |  
47113 generated, thus making the signal delivery resources available for delivery of subsequent |  
47114 periodic timer expiration signals. But, in general, this cannot be assured—processing of periodic |  
47115 timer signals may “overrun”; that is, subsequent periodic timer expirations may occur before the |  
47116 currently pending signal has been delivered.

47117 Also, for signals, according to the POSIX.1-1990 standard, if subsequent occurrences of a |  
47118 pending signal are generated, it is implementation-defined whether a signal is delivered for each |  
47119 occurrence. This is not adequate for some realtime applications. So a mechanism is required to |  
47120 allow applications to detect how many timer expirations were delayed without requiring an |  
47121 indefinite amount of system resources to store the delayed expirations.

47122 The specified facilities provide for an overrun count. The overrun count is defined as the number |  
47123 of extra timer expirations that occurred between the time a timer expiration signal is generated |  
47124 and the time the signal is delivered. The signal-catching function, if it is concerned with |  
47125 overruns, can retrieve this count on entry. With this method, a periodic timer only needs one |  
47126 “signal queuing resource” that can be allocated at the time of the *timer\_create()* function call.

47127 A function is defined to retrieve the overrun count so that an application need not allocate static |  
47128 storage to contain the count, and an implementation need not update this storage |  
47129 asynchronously on timer expirations. But, for some high-frequency periodic applications, the |  
47130 overhead of an additional system call on each timer expiration may be prohibitive. The |  
47131 functions, as defined, permit an implementation to maintain the overrun count in user space, |  
47132 associated with the *timerid*. The *timer\_getoverrun()* function can then be implemented as a macro |  
47133 that uses the *timerid* argument (which may just be a pointer to a user space structure containing |  
47134 the counter) to locate the overrun count with no system call overhead. Other implementations, |  
47135 less concerned with this class of applications, can avoid the asynchronous update of user space |  
47136 by maintaining the count in a system structure at the cost of the extra system call to obtain it.

47137 **Timer Expiration Signal Parameters**

47138 The Realtime Signals Extension option supports an application-specific datum that is delivered  
47139 to the extended signal handler. This value is explicitly specified by the application, along with  
47140 the signal number to be delivered, in a **sigevent** structure. The type of the application-defined  
47141 value can be either an integer constant or a pointer. This explicit specification of the value, as  
47142 opposed to always sending the timer ID, was selected based on existing practice.

47143 It is common practice for realtime applications (on non-POSIX systems or realtime extended  
47144 POSIX systems) to use the parameters of event handlers as the case label of a switch statement  
47145 or as a pointer to an application-defined data structure. Since *timer\_ids* are dynamically allocated  
47146 by the *timer\_create()* function, they can be used for neither of these functions without additional  
47147 application overhead in the signal handler; for example, to search an array of saved timer IDs to  
47148 associate the ID with a constant or application data structure.

47149 **FUTURE DIRECTIONS**

47150 None.

47151 **SEE ALSO**

47152 *clock\_getres()*, *timer\_delete()*, *timer\_getoverrun()*, the Base Definitions volume of  
47153 IEEE Std 1003.1-200x, <**time.h**>

47154 **CHANGE HISTORY**

47155 First released in Issue 5. Included for alignment with the POSIX Realtime Extension.

47156 **Issue 6**

47157 The *timer\_create()* function is marked as part of the Timers option.

47158 The [ENOSYS] error condition has been removed as stubs need not be provided if an  
47159 implementation does not support the Timers option.

47160 CPU-time clocks are added for alignment with IEEE Std 1003.1d-1999.

47161 The DESCRIPTION is updated for alignment with IEEE Std 1003.1j-2000 by adding the  
47162 requirement for the CLOCK\_MONOTONIC clock under the Monotonic Clock option.

47163 The **restrict** keyword is added to the *timer\_create()* prototype for alignment with the  
47164 ISO/IEC 9899:1999 standard.

47165 **NAME**

47166 timer\_delete — delete a per-process timer (**REALTIME**)

47167 **SYNOPSIS**

```
47168 TMR #include <time.h>
```

```
47169 int timer_delete(timer_t timerid);
```

47170

47171 **DESCRIPTION**

47172 The *timer\_delete()* function deletes the specified timer, *timerid*, previously created by the  
47173 *timer\_create()* function. If the timer is armed when *timer\_delete()* is called, the behavior shall be  
47174 as if the timer is automatically disarmed before removal. The disposition of pending signals for  
47175 the deleted timer is unspecified.

47176 **RETURN VALUE**

47177 If successful, the *timer\_delete()* function shall return a value of zero. Otherwise, the function shall  
47178 return a value of  $-1$  and set *errno* to indicate the error.

47179 **ERRORS**

47180 The *timer\_delete()* function shall fail if:

47181 [EINVAL] The timer ID specified by *timerid* is not a valid timer ID.

47182 **EXAMPLES**

47183 None.

47184 **APPLICATION USAGE**

47185 None.

47186 **RATIONALE**

47187 None.

47188 **FUTURE DIRECTIONS**

47189 None.

47190 **SEE ALSO**

47191 *timer\_create()*, the Base Definitions volume of IEEE Std 1003.1-200x, <**time.h**>

47192 **CHANGE HISTORY**

47193 First released in Issue 5. Included for alignment with the POSIX Realtime Extension.

47194 **Issue 6**

47195 The *timer\_delete()* function is marked as part of the Timers option.

47196 The [ENOSYS] error condition has been removed as stubs need not be provided if an  
47197 implementation does not support the Timers option.

## 47198 NAME

47199 timer\_getoverrun, timer\_gettime, timer\_settime — per-process timers (REALTIME)

## 47200 SYNOPSIS

```
47201 TMR    #include <time.h>
47202
47202     int timer_getoverrun(timer_t timerid);
47203     int timer_gettime(timer_t timerid, struct itimerspec *value);
47204     int timer_settime(timer_t timerid, int flags,
47205                      const struct itimerspec *restrict value,
47206                      struct itimerspec *restrict ovalue);
47207
```

## 47208 DESCRIPTION

47209 The *timer\_gettime()* function shall store the amount of time until the specified timer, *timerid*, |  
 47210 expires and the reload value of the timer into the space pointed to by the *value* argument. The |  
 47211 *it\_value* member of this structure shall contain the amount of time before the timer expires, or |  
 47212 zero if the timer is disarmed. This value is returned as the interval until timer expiration, even if |  
 47213 the timer was armed with absolute time. The *it\_interval* member of *value* shall contain the reload |  
 47214 value last set by *timer\_settime()*.

47215 The *timer\_settime()* function shall set the time until the next expiration of the timer specified by |  
 47216 *timerid* from the *it\_value* member of the *value* argument and arms the timer if the *it\_value* |  
 47217 member of *value* is non-zero. If the specified timer was already armed when *timer\_settime()* is |  
 47218 called, this call shall reset the time until next expiration to the *value* specified. If the *it\_value* |  
 47219 member of *value* is zero, the timer shall be disarmed. The effect of disarming or resetting a timer |  
 47220 with pending expiration notifications is unspecified.

47221 If the flag `TIMER_ABSTIME` is not set in the argument *flags*, *timer\_settime()* shall behave as if the |  
 47222 time until next expiration is set to be equal to the interval specified by the *it\_value* member of |  
 47223 *value*. That is, the timer shall expire in *it\_value* nanoseconds from when the call is made. If the |  
 47224 flag `TIMER_ABSTIME` is set in the argument *flags*, *timer\_settime()* shall behave as if the time |  
 47225 until next expiration is set to be equal to the difference between the absolute time specified by |  
 47226 the *it\_value* member of *value* and the current value of the clock associated with *timerid*. That is, |  
 47227 the timer shall expire when the clock reaches the value specified by the *it\_value* member of *value*. |  
 47228 If the specified time has already passed, the function shall succeed and the expiration |  
 47229 notification shall be made.

47230 The reload value of the timer shall be set to the value specified by the *it\_interval* member of |  
 47231 *value*. When a timer is armed with a non-zero *it\_interval*, a periodic (or repetitive) timer is |  
 47232 specified.

47233 Time values that are between two consecutive non-negative integer multiples of the resolution |  
 47234 of the specified timer shall be rounded up to the larger multiple of the resolution. Quantization |  
 47235 error shall not cause the timer to expire earlier than the rounded time value.

47236 If the argument *ovalue* is not NULL, the function *timer\_settime()* shall store, in the location |  
 47237 referenced by *ovalue*, a value representing the previous amount of time before the timer would |  
 47238 have expired, or zero if the timer was disarmed, together with the previous timer reload value. |  
 47239 Timers shall not expire before their scheduled time.

47240 Only a single signal shall be queued to the process for a given timer at any point in time. When a |  
 47241 timer for which a signal is still pending expires, no signal shall be queued, and a timer overrun |  
 47242 shall occur. When a timer expiration signal is delivered to or accepted by a process, if the |  
 47243 implementation supports the Realtime Signals Extension, the *timer\_getoverrun()* function shall |  
 47244 return the timer expiration overrun count for the specified timer. The overrun count returned |  
 47245 contains the number of extra timer expirations that occurred between the time the signal was

47246 generated (queued) and when it was delivered or accepted, up to but not including an  
47247 implementation-defined maximum of {DELAYTIMER\_MAX}. If the number of such extra  
47248 expirations is greater than or equal to {DELAYTIMER\_MAX}, then the overrun count shall be set  
47249 to {DELAYTIMER\_MAX}. The value returned by *timer\_getoverrun()* shall apply to the most  
47250 recent expiration signal delivery or acceptance for the timer. If no expiration signal has been  
47251 delivered for the timer, or if the Realtime Signals Extension is not supported, the return value of  
47252 *timer\_getoverrun()* is unspecified.

#### 47253 RETURN VALUE

47254 If the *timer\_getoverrun()* function succeeds, it shall return the timer expiration overrun count as  
47255 explained above.

47256 If the *timer\_gettime()* or *timer\_settime()* functions succeed, a value of 0 shall be returned.

47257 If an error occurs for any of these functions, the value -1 shall be returned, and *errno* set to  
47258 indicate the error.

#### 47259 ERRORS

47260 The *timer\_getoverrun()*, *timer\_gettime()*, and *timer\_settime()* functions shall if:

47261 [EINVAL] The *timerid* argument does not correspond to an ID returned by *timer\_create()*  
47262 but not yet deleted by *timer\_delete()*.

47263 The *timer\_settime()* function shall fail if:

47264 [EINVAL] A *value* structure specified a nanosecond value less than zero or greater than  
47265 or equal to 1,000 million, and the *it\_value* member of that structure did not  
47266 specify zero seconds and nanoseconds.

#### 47267 EXAMPLES

47268 None.

#### 47269 APPLICATION USAGE

47270 None.

#### 47271 RATIONALE

47272 Practical clocks tick at a finite rate, with rates of 100 Hertz and 1,000 Hertz being common. The  
47273 inverse of this tick rate is the clock resolution, also called the clock granularity, which in either  
47274 case is expressed as a time duration, being 10 milliseconds and 1 millisecond respectively for  
47275 these common rates. The granularity of practical clocks implies that if one reads a given clock  
47276 twice in rapid succession, one may get the same time value twice; and that timers must wait for  
47277 the next clock tick after the theoretical expiration time, to ensure that a timer never returns too  
47278 soon. Note also that the granularity of the clock may be significantly coarser than the resolution  
47279 of the data format used to set and get time and interval values. Also note that some  
47280 implementations may choose to adjust time and/or interval values to exactly match the ticks of  
47281 the underlying clock.

47282 This volume of IEEE Std 1003.1-200x defines functions that allow an application to determine the  
47283 implementation-supported resolution for the clocks and requires an implementation to  
47284 document the resolution supported for timers and *nanosleep()* if they differ from the supported  
47285 clock resolution. This is more of a procurement issue than a runtime application issue.

#### 47286 FUTURE DIRECTIONS

47287 None.

47288 **SEE ALSO**

47289 *clock\_getres()*, *timer\_create()*, the Base Definitions volume of IEEE Std 1003.1-200x, <**time.h**>

47290 **CHANGE HISTORY**

47291 First released in Issue 5. Included for alignment with the POSIX Realtime Extension.

47292 **Issue 6**

47293 The *timer\_getoverrun()*, *timer\_gettime()*, and *timer\_settime()* functions are marked as part of the  
47294 Timers option.

47295 The [ENOSYS] error condition has been removed as stubs need not be provided if an  
47296 implementation does not support the Timers option.

47297 The [EINVAL] error condition is updated to include the following: “and the *it\_value* member of  
47298 that structure did not specify zero seconds and nanoseconds.” This change is for IEEE PASC  
47299 Interpretation 1003.1 #89.

47300 The DESCRIPTION for *timer\_getoverrun()* is updated to clarify that “If no expiration signal has  
47301 been delivered for the timer, or if the Realtime Signals Extension is not supported, the return  
47302 value of *timer\_getoverrun()* is unspecified”.

47303 The **restrict** keyword is added to the *timer\_settime()* prototype for alignment with the  
47304 ISO/IEC 9899:1999 standard.

47305 **NAME**

47306 times — get process and waited-for child process times

47307 **SYNOPSIS**

47308 #include &lt;sys/times.h&gt;

47309 clock\_t times(struct tms \*buffer);

47310 **DESCRIPTION**47311 The *times()* function shall fill the **tms** structure pointed to by *buffer* with time-accounting  
47312 information. The **tms** structure is defined in <sys/times.h>.

47313 All times are measured in terms of the number of clock ticks used.

47314 The times of a terminated child process shall be included in the *tms\_cutime* and *tms\_cstime*  
47315 elements of the parent when *wait()* or *waitpid()* returns the process ID of this terminated child. If  
47316 a child process has not waited for its children, their times shall not be included in its times.47317 • The *tms\_utime* structure member is the CPU time charged for the execution of user  
47318 instructions of the calling process.47319 • The *tms\_stime* structure member is the CPU time charged for execution by the system on  
47320 behalf of the calling process.47321 • The *tms\_cutime* structure member is the sum of the *tms\_utime* and *tms\_cutime* times of the  
47322 child processes.47323 • The *tms\_cstime* structure member is the sum of the *tms\_stime* and *tms\_cstime* times of the child  
47324 processes.47325 **RETURN VALUE**47326 Upon successful completion, *times()* shall return the elapsed real time, in clock ticks, since an  
47327 arbitrary point in the past (for example, system start-up time). This point does not change from  
47328 one invocation of *times()* within the process to another. The return value may overflow the  
47329 possible range of type **clock\_t**. If *times()* fails, (**clock\_t**)-1 shall be returned and *errno* set to  
47330 indicate the error.47331 **ERRORS**

47332 No errors are defined.

47333 **EXAMPLES**47334 **Timing a Database Lookup**47335 The following example defines two functions, *start\_clock()* and *end\_clock()*, that are used to time  
47336 a lookup. It also defines variables of type **clock\_t** and **tms** to measure the duration of  
47337 transactions. The *start\_clock()* function saves the beginning times given by the *times()* function.  
47338 The *end\_clock()* function gets the ending times and prints the difference between the two times.47339 #include <sys/times.h>  
47340 #include <stdio.h>  
47341 ...  
47342 void start\_clock(void);  
47343 void end\_clock(char \*msg);  
47344 ...  
47345 static clock\_t st\_time;  
47346 static clock\_t en\_time;  
47347 static struct tms st\_cpu;  
47348 static struct tms en\_cpu;



```

47349     ...
47350     void
47351     start_clock()
47352     {
47353         st_time = times(&st_cpu);
47354     }
47355     void
47356     end_clock(char *msg)
47357     {
47358         en_time = times(&en_cpu);
47359
47359         printf(msg);
47360         printf("Real Time: %ld, User Time %ld, System Time %ld\n",
47361             en_time - st_time,
47362             en_cpu.tms_utime - st_cpu.tms_utime,
47363             en_cpu.tms_stime - st_cpu.tms_stime);
47364     }

```

**47365 APPLICATION USAGE**

47366 Applications should use `sysconf(_SC_CLK_TCK)` to determine the number of clock ticks per  
47367 second as it may vary from system to system.

**47368 RATIONALE**

47369 The accuracy of the times reported is intentionally left unspecified to allow implementations  
47370 flexibility in design, from uniprocessor to multi-processor networks.

47371 The inclusion of times of child processes is recursive, so that a parent process may collect the  
47372 total times of all of its descendants. But the times of a child are only added to those of its parent  
47373 when its parent successfully waits on the child. Thus, it is not guaranteed that a parent process  
47374 can always see the total times of all its descendants; see also the discussion of the term *realtime* in  
47375 *alarm()*.

47376 If the type `clock_t` is defined to be a signed 32-bit integer, it overflows in somewhat more than a  
47377 year if there are 60 clock ticks per second, or less than a year if there are 100. There are individual  
47378 systems that run continuously for longer than that. This volume of IEEE Std 1003.1-200x permits  
47379 an implementation to make the reference point for the returned value be the start-up time of the  
47380 process, rather than system start-up time.

47381 The term *charge* in this context has nothing to do with billing for services. The operating system  
47382 accounts for time used in this way. That information must be correct, regardless of how that  
47383 information is used.

**47384 FUTURE DIRECTIONS**

47385 None.

**47386 SEE ALSO**

47387 `exec`, `fork()`, `sysconf()`, `time()`, `wait()`, the Base Definitions volume of IEEE Std 1003.1-200x,  
47388 `<sys/times.h>`

**47389 CHANGE HISTORY**

47390 First released in Issue 1. Derived from Issue 1 of the SVID.

47391 **NAME**

47392            timezone — difference from UTC and local standard time

47393 **SYNOPSIS**

47394 xSI        #include <time.h>

47395            extern long timezone;

47396

47397 **DESCRIPTION**

47398            Refer to *tzset()*.

47399 **NAME**

47400 tmpfile — create a temporary file

47401 **SYNOPSIS**

47402 #include &lt;stdio.h&gt;

47403 FILE \*tmpfile(void);

47404 **DESCRIPTION**

47405 CX The functionality described on this reference page is aligned with the ISO C standard. Any  
 47406 conflict between the requirements described here and the ISO C standard is unintentional. This  
 47407 volume of IEEE Std 1003.1-200x defers to the ISO C standard.

47408 The *tmpfile()* function shall create a temporary file and open a corresponding stream. The file  
 47409 shall be automatically deleted when all references to the file are closed. The file is opened as in  
 47410 *fopen()* for update (*w+*).

47411 CX In some implementations, a permanent file may be left behind if the process calling *tmpfile()* is  
 47412 killed while it is processing a call to *tmpfile()*.

47413 An error message may be written to standard error if the stream cannot be opened.

47414 **RETURN VALUE**

47415 Upon successful completion, *tmpfile()* shall return a pointer to the stream of the file that is  
 47416 CX created. Otherwise, it shall return a null pointer and set *errno* to indicate the error.

47417 **ERRORS**47418 The *tmpfile()* function shall fail if:

47419 CX [EINTR] A signal was caught during *tmpfile()*.

47420 CX [EMFILE] {OPEN\_MAX} file descriptors are currently open in the calling process.

47421 CX [ENFILE] The maximum allowable number of files is currently open in the system.

47422 CX [ENOSPC] The directory or file system which would contain the new file cannot be  
 47423 expanded.

47424 CX [EOVERFLOW] The file is a regular file and the size of the file cannot be represented correctly  
 47425 in an object of type *off\_t*.

47426 The *tmpfile()* function may fail if:

47427 CX [EMFILE] {FOPEN\_MAX} streams are currently open in the calling process.

47428 CX [ENOMEM] Insufficient storage space is available.

47429 **EXAMPLES**47430 **Creating a Temporary File**

47431 The following example creates a temporary file for update, and returns a pointer to a stream for  
 47432 the created file in the *fp* variable.

47433 #include &lt;stdio.h&gt;

47434 ...

47435 FILE \*fp;

47436 fp = tmpfile ();

47437 **APPLICATION USAGE**

47438 It should be possible to open at least {TMP\_MAX} temporary files during the lifetime of the  
47439 program (this limit may be shared with *tmpnam()*) and there should be no limit on the number  
47440 simultaneously open other than this limit and any limit on the number of open files  
47441 ({FOPEN\_MAX}).

47442 **RATIONALE**

47443 None.

47444 **FUTURE DIRECTIONS**

47445 None.

47446 **SEE ALSO**

47447 *fopen()*, *tmpnam()*, *unlink()*, the Base Definitions volume of IEEE Std 1003.1-200x, <stdio.h>

47448 **CHANGE HISTORY**

47449 First released in Issue 1. Derived from Issue 1 of the SVID.

47450 **Issue 5**

47451 Large File Summit extensions are added.

47452 The last two paragraphs of the DESCRIPTION were included as APPLICATION USAGE notes  
47453 in previous issues.

47454 **Issue 6**

47455 Extensions beyond the ISO C standard are now marked.

47456 The following new requirements on POSIX implementations derive from alignment with the  
47457 Single UNIX Specification:

- 47458 • In the ERRORS section, the [Eoverflow] condition is added. This change is to support  
47459 large files.
- 47460 • The [EMFILE] optional error condition is added.

47461 The APPLICATION USAGE section is added for alignment with the ISO/IEC 9899:1999  
47462 standard.

47463 **NAME**

47464 tmpnam — create a name for a temporary file

47465 **SYNOPSIS**

47466 #include &lt;stdio.h&gt;

47467 char \*tmpnam(char \*s);

47468 **DESCRIPTION**

47469 cx The functionality described on this reference page is aligned with the ISO C standard. Any  
 47470 conflict between the requirements described here and the ISO C standard is unintentional. This  
 47471 volume of IEEE Std 1003.1-200x defers to the ISO C standard.

47472 The *tmpnam()* function shall generate a string that is a valid filename and that is not the same as  
 47473 the name of an existing file. The function is potentially capable of generating {TMP\_MAX}  
 47474 different strings, but any or all of them may already be in use by existing files and thus not be  
 47475 suitable return values.

47476 The *tmpnam()* function generates a different string each time it is called from the same process,  
 47477 up to {TMP\_MAX} times. If it is called more than {TMP\_MAX} times, the behavior is  
 47478 implementation-defined.

47479 The implementation shall behave as if no function defined in this volume of  
 47480 IEEE Std 1003.1-200x calls *tmpnam()*.

47481 cx If the application uses any of the functions guaranteed to be available if either  
 47482 \_POSIX\_THREAD\_SAFE\_FUNCTIONS or \_POSIX\_THREADS is defined, the application shall  
 47483 ensure that the *tmpnam()* function is called with a non-NULL parameter.

47484 **RETURN VALUE**

47485 Upon successful completion, *tmpnam()* shall return a pointer to a string. If no suitable string can  
 47486 be generated, the *tmpnam()* function shall return a null pointer.

47487 If the argument *s* is a null pointer, *tmpnam()* shall leave its result in an internal static object and  
 47488 return a pointer to that object. Subsequent calls to *tmpnam()* may modify the same object. If the  
 47489 argument *s* is not a null pointer, it is presumed to point to an array of at least L\_tmpnam **chars**;  
 47490 *tmpnam()* shall write its result in that array and shall return the argument as its value.

47491 **ERRORS**

47492 No errors are defined.

47493 **EXAMPLES**47494 **Generating a Filename**47495 The following example generates a unique filename and stores it in the array pointed to by *ptr*.

47496 #include &lt;stdio.h&gt;

47497 ...

47498 char filename[L\_tmpnam+1];

47499 char \*ptr;

47500 ptr = tmpnam(filename);

47501 **APPLICATION USAGE**

47502 This function only creates filenames. It is the application's responsibility to create and remove  
 47503 the files.

47504 Between the time a pathname is created and the file is opened, it is possible for some other  
 47505 process to create a file with the same name. Applications may find *tmpfile()* more useful.

47506 **RATIONALE**

47507           None.

47508 **FUTURE DIRECTIONS**

47509           None.

47510 **SEE ALSO**47511           *fopen()*, *open()*, *tmpnam()*, *tmpfile()*, *unlink()*, the Base Definitions volume of  
47512           IEEE Std 1003.1-200x, <**stdio.h**>47513 **CHANGE HISTORY**

47514           First released in Issue 1. Derived from Issue 1 of the SVID.

47515 **Issue 5**

47516           The DESCRIPTION is updated for alignment with the POSIX Threads Extension.

47517 **Issue 6**

47518           Extensions beyond the ISO C standard are now marked.

47519           The DESCRIPTION is updated to avoid use of the term “must” for application requirements.

47520           The DESCRIPTION is expanded for alignment with the ISO/IEC 9899:1999 standard.

47521 **NAME**

47522           toascii — translate integer to a 7-bit ASCII character

47523 **SYNOPSIS**

47524 xSI       #include &lt;ctype.h&gt;

47525           int toascii(int c);

47526

47527 **DESCRIPTION**47528           The *toascii()* function shall convert its argument into a 7-bit ASCII character.47529 **RETURN VALUE**47530           The *toascii()* function shall return the value (*c* &0x7f).47531 **ERRORS**

47532           No errors are returned.

47533 **EXAMPLES**

47534           None.

47535 **APPLICATION USAGE**

47536           None.

47537 **RATIONALE**

47538           None.

47539 **FUTURE DIRECTIONS**

47540           None.

47541 **SEE ALSO**47542           *isascii()*, the Base Definitions volume of IEEE Std 1003.1-200x, <ctype.h>47543 **CHANGE HISTORY**

47544           First released in Issue 1. Derived from Issue 1 of the SVID.

47545 **NAME**

47546           tolower — transliterate uppercase characters to lowercase

47547 **SYNOPSIS**

47548           #include &lt;ctype.h&gt;

47549           int tolower(int c);

47550 **DESCRIPTION**

47551 cx       The functionality described on this reference page is aligned with the ISO C standard. Any  
47552       conflict between the requirements described here and the ISO C standard is unintentional. This  
47553       volume of IEEE Std 1003.1-200x defers to the ISO C standard.

47554       The *tolower()* function has as a domain a type **int**, the value of which is representable as an  
47555       **unsigned char** or the value of EOF. If the argument has any other value, the behavior is  
47556       undefined. If the argument of *tolower()* represents an uppercase letter, and there exists a  
47557 cx       corresponding lowercase letter (as defined by character type information in the program locale  
47558       category *LC\_CTYPE*), the result shall be the corresponding lowercase letter. All other arguments  
47559       in the domain are returned unchanged. |

47560 **RETURN VALUE**

47561       Upon successful completion, *tolower()* shall return the lowercase letter corresponding to the  
47562       argument passed; otherwise, it shall return the argument unchanged.

47563 **ERRORS**

47564       No errors are defined.

47565 **EXAMPLES**

47566       None.

47567 **APPLICATION USAGE**

47568       None.

47569 **RATIONALE**

47570       None.

47571 **FUTURE DIRECTIONS**

47572       None.

47573 **SEE ALSO**

47574       *setlocale()*, the Base Definitions volume of IEEE Std 1003.1-200x, <ctype.h>, the Base Definitions  
47575       volume of IEEE Std 1003.1-200x, Chapter 7, Locale

47576 **CHANGE HISTORY**

47577       First released in Issue 1. Derived from Issue 1 of the SVID.

47578 **Issue 6**

47579       Extensions beyond the ISO C standard are now marked.



47580 **NAME**

47581           toupper — transliterate lowercase characters to uppercase

47582 **SYNOPSIS**

47583           #include &lt;ctype.h&gt;

47584           int toupper(int c);

47585 **DESCRIPTION**

47586 cx       The functionality described on this reference page is aligned with the ISO C standard. Any  
47587       conflict between the requirements described here and the ISO C standard is unintentional. This  
47588       volume of IEEE Std 1003.1-200x defers to the ISO C standard.

47589       The *toupper()* function has as a domain a type **int**, the value of which is representable as an  
47590       **unsigned char** or the value of EOF. If the argument has any other value, the behavior is  
47591       undefined. If the argument of *toupper()* represents a lowercase letter, and there exists a  
47592 cx       corresponding uppercase letter (as defined by character type information in the program locale  
47593       category *LC\_CTYPE*), the result shall be the corresponding uppercase letter. All other arguments  
47594       in the domain are returned unchanged. |

47595 **RETURN VALUE**

47596       Upon successful completion, *toupper()* shall return the uppercase letter corresponding to the  
47597       argument passed.

47598 **ERRORS**

47599       No errors are defined.

47600 **EXAMPLES**

47601       None.

47602 **APPLICATION USAGE**

47603       None.

47604 **RATIONALE**

47605       None.

47606 **FUTURE DIRECTIONS**

47607       None.

47608 **SEE ALSO**

47609       *setlocale()*, the Base Definitions volume of IEEE Std 1003.1-200x, <ctype.h>, the Base Definitions  
47610       volume of IEEE Std 1003.1-200x, Chapter 7, Locale

47611 **CHANGE HISTORY**

47612       First released in Issue 1. Derived from Issue 1 of the SVID.

47613 **Issue 6**

47614       Extensions beyond the ISO C standard are now marked.

47615 **NAME**

47616 towctrans — wide-character transliteration

47617 **SYNOPSIS**

47618 #include &lt;wctype.h&gt;

47619 wint\_t towctrans(wint\_t *wc*, wctrans\_t *desc*);47620 **DESCRIPTION**

47621 CX The functionality described on this reference page is aligned with the ISO C standard. Any  
 47622 conflict between the requirements described here and the ISO C standard is unintentional. This  
 47623 volume of IEEE Std 1003.1-200x defers to the ISO C standard.

47624 The *towctrans()* function shall transliterate the wide-character code *wc* using the mapping  
 47625 described by *desc*. The current setting of the *LC\_CTYPE* category should be the same as during  
 47626 CX the call to *wctrans()* that returned the value *desc*. If the value of *desc* is invalid (that is, not  
 47627 obtained by a call to *wctrans()* or *desc* is invalidated by a subsequent call to *setlocale()* that has  
 47628 affected category *LC\_CTYPE*), the result is unspecified.

47629 An application wishing to check for error situations should set *errno* to 0 before calling  
 47630 *towctrans()*. If *errno* is non-zero on return, an error has occurred.

47631 **RETURN VALUE**

47632 If successful, the *towctrans()* function shall return the mapped value of *wc* using the mapping  
 47633 described by *desc*. Otherwise, it shall return *wc* unchanged.

47634 **ERRORS**47635 The *towctrans()* function may fail if:47636 CX [EINVAL] *desc* contains an invalid transliteration descriptor.47637 **EXAMPLES**

47638 None.

47639 **APPLICATION USAGE**

47640 The strings "tolower" and "toupper" are reserved for the standard mapping names. In the  
 47641 table below, the functions in the left column are equivalent to the functions in the right column.

47642 tolower(*wc*) towctrans(*wc*, wctrans("tolower"))47643 toupper(*wc*) towctrans(*wc*, wctrans("toupper"))47644 **RATIONALE**

47645 None.

47646 **FUTURE DIRECTIONS**

47647 None.

47648 **SEE ALSO**

47649 *tolower()*, *toupper()*, *wctrans()*, the Base Definitions volume of IEEE Std 1003.1-200x,  
 47650 <wctype.h>

47651 **CHANGE HISTORY**

47652 First released in Issue 5. Derived from ISO/IEC 9899:1990/Amendment 1:1995 (E).

47653 **Issue 6**

47654 Extensions beyond the ISO C standard are now marked.

47655 **NAME**

47656 tolower — transliterate uppercase wide-character code to lowercase

47657 **SYNOPSIS**

47658 #include &lt;wctype.h&gt;

47659 wint\_t tolower(wint\_t wc);

47660 **DESCRIPTION**

47661 cx The functionality described on this reference page is aligned with the ISO C standard. Any  
47662 conflict between the requirements described here and the ISO C standard is unintentional. This  
47663 volume of IEEE Std 1003.1-200x defers to the ISO C standard.

47664 The *tolower()* function has as a domain a type **wint\_t**, the value of which the application shall  
47665 ensure is a character representable as a **wchar\_t**, and a wide-character code corresponding to a  
47666 valid character in the current locale or the value of WEOF. If the argument has any other value,  
47667 the behavior is undefined. If the argument of *tolower()* represents an uppercase wide-character  
47668 code, and there exists a corresponding lowercase wide-character code (as defined by character  
47669 type information in the program locale category *LC\_CTYPE*), the result shall be the  
47670 corresponding lowercase wide-character code. All other arguments in the domain are returned  
47671 unchanged.

47672 **RETURN VALUE**

47673 Upon successful completion, *tolower()* shall return the lowercase letter corresponding to the  
47674 argument passed; otherwise, it shall return the argument unchanged.

47675 **ERRORS**

47676 No errors are defined.

47677 **EXAMPLES**

47678 None.

47679 **APPLICATION USAGE**

47680 None.

47681 **RATIONALE**

47682 None.

47683 **FUTURE DIRECTIONS**

47684 None.

47685 **SEE ALSO**

47686 *setlocale()*, the Base Definitions volume of IEEE Std 1003.1-200x, <wctype.h>, <wchar.h>, the  
47687 Base Definitions volume of IEEE Std 1003.1-200x, Chapter 7, Locale

47688 **CHANGE HISTORY**

47689 First released in Issue 4.

47690 **Issue 5**

47691 The following change has been made in this issue for alignment with  
47692 ISO/IEC 9899:1990/Amendment 1:1995 (E):

- 47693 • The SYNOPSIS has been changed to indicate that this function and associated data types are  
47694 now made visible by inclusion of the <wctype.h> header rather than <wchar.h>.

47695 **Issue 6**

47696 The DESCRIPTION is updated to avoid use of the term “must” for application requirements.

47697 **NAME**

47698 towupper — transliterate lowercase wide-character code to uppercase

47699 **SYNOPSIS**

47700 #include &lt;wctype.h&gt;

47701 wint\_t towupper(wint\_t wc);

47702 **DESCRIPTION**

47703 cx The functionality described on this reference page is aligned with the ISO C standard. Any  
47704 conflict between the requirements described here and the ISO C standard is unintentional. This  
47705 volume of IEEE Std 1003.1-200x defers to the ISO C standard.

47706 The *towupper()* function has as a domain a type **wint\_t**, the value of which the application shall  
47707 ensure is a character representable as a **wchar\_t**, and a wide-character code corresponding to a  
47708 valid character in the current locale or the value of WEOF. If the argument has any other value,  
47709 the behavior is undefined. If the argument of *towupper()* represents a lowercase wide-character  
47710 code, and there exists a corresponding uppercase wide-character code (as defined by character  
47711 type information in the program locale category *LC\_CTYPE*), the result shall be the  
47712 corresponding uppercase wide-character code. All other arguments in the domain are returned  
47713 unchanged.

47714 **RETURN VALUE**

47715 Upon successful completion, *towupper()* shall return the uppercase letter corresponding to the  
47716 argument passed. Otherwise, it shall return the argument unchanged.

47717 **ERRORS**

47718 No errors are defined.

47719 **EXAMPLES**

47720 None.

47721 **APPLICATION USAGE**

47722 None.

47723 **RATIONALE**

47724 None.

47725 **FUTURE DIRECTIONS**

47726 None.

47727 **SEE ALSO**

47728 *setlocale()*, the Base Definitions volume of IEEE Std 1003.1-200x, <wctype.h>, <wchar.h>, the  
47729 Base Definitions volume of IEEE Std 1003.1-200x, Chapter 7, Locale

47730 **CHANGE HISTORY**

47731 First released in Issue 4.

47732 **Issue 5**

47733 The following change has been made in this issue for alignment with  
47734 ISO/IEC 9899:1990/Amendment 1:1995 (E):

- 47735 • The SYNOPSIS has been changed to indicate that this function and associated data types are  
47736 now made visible by inclusion of the <wctype.h> header rather than <wchar.h>.

47737 **Issue 6**

47738 The DESCRIPTION is updated to avoid use of the term “must” for application requirements.

47739 **NAME**

47740 trunc, truncf, trunc1 — round to truncated integer value

47741 **SYNOPSIS**

47742 #include &lt;math.h&gt;

47743 double trunc(double x);

47744 float truncf(float x);

47745 long double trunc1(long double x);

47746 **DESCRIPTION**47747 **CX** The functionality described on this reference page is aligned with the ISO C standard. Any  
47748 conflict between the requirements described here and the ISO C standard is unintentional. This  
47749 volume of IEEE Std 1003.1-200x defers to the ISO C standard.47750 These functions shall round their argument to the integer value, in floating format, nearest to but  
47751 no larger in magnitude than the argument.47752 **RETURN VALUE**

47753 Upon successful completion, these functions shall return the truncated integer value.

47754 **MX** If  $x$  is NaN, a NaN shall be returned.47755 If  $x$  is  $\pm 0$ , or  $\pm \text{Inf}$ ,  $x$  shall be returned.47756 **ERRORS**

47757 No errors are defined.

47758 **EXAMPLES**

47759 None.

47760 **APPLICATION USAGE**

47761 None.

47762 **RATIONALE**

47763 None.

47764 **FUTURE DIRECTIONS**

47765 None.

47766 **SEE ALSO**

47767 The Base Definitions volume of IEEE Std 1003.1-200x, &lt;math.h&gt;

47768 **CHANGE HISTORY**

47769 First released in Issue 6. Derived from the ISO/IEC 9899:1999 standard.

47770 **NAME**

47771 truncate — truncate a file to a specified length

47772 **SYNOPSIS**47773 XSI `#include <unistd.h>`47774 `int truncate(const char *path, off_t length);`

47775

47776 **DESCRIPTION**47777 The `truncate()` function shall cause the regular file named by *path* to have a size which shall be  
47778 equal to *length* bytes.47779 If the file previously was larger than *length*, the extra data is discarded. If the file was previously  
47780 shorter than *length*, its size is increased, and the extended area appears as if it were zero-filled.

47781 The application shall ensure that the process has write permission for the file.

47782 If the request would cause the file size to exceed the soft file size limit for the process, the  
47783 request shall fail and the implementation shall generate the SIGXFSZ signal for the process.47784 This function shall not modify the file offset for any open file descriptions associated with the  
47785 file. Upon successful completion, if the file size is changed, this function shall mark for update  
47786 the *st\_ctime* and *st\_mtime* fields of the file, and the S\_ISUID and S\_ISGID bits of the file mode  
47787 may be cleared.47788 **RETURN VALUE**47789 Upon successful completion, `truncate()` shall return 0. Otherwise, `-1` shall be returned, and *errno*  
47790 set to indicate the error.47791 **ERRORS**47792 The `truncate()` function shall fail if:

47793 [EINTR] A signal was caught during execution.

47794 [EINVAL] The *length* argument was less than 0.

47795 [EFBIG] or [EINVAL]

47796 The *length* argument was greater than the maximum file size.

47797 [EIO] An I/O error occurred while reading from or writing to a file system.

47798 [EACCES] A component of the path prefix denies search permission, or write permission  
47799 is denied on the file.

47800 [EISDIR] The named file is a directory.

47801 [ELOOP] A loop exists in symbolic links encountered during resolution of the *path*  
47802 argument.

47803 [ENAMETOOLONG]

47804 The length of the *path* argument exceeds {PATH\_MAX} or a pathname |  
47805 component is longer than {NAME\_MAX}. |47806 [ENOENT] A component of *path* does not name an existing file or *path* is an empty string.47807 [ENOTDIR] A component of the path prefix of *path* is not a directory.

47808 [EROFS] The named file resides on a read-only file system.

47809 The `truncate()` function may fail if:

- 47810 [ELOOP] More than {SYMLOOP\_MAX} symbolic links were encountered during  
47811 resolution of the *path* argument.
- 47812 [ENAMETOOLONG]  
47813 Pathname resolution of a symbolic link produced an intermediate result |  
47814 whose length exceeds {PATH\_MAX}.
- 47815 **EXAMPLES**  
47816 None.
- 47817 **APPLICATION USAGE**  
47818 None.
- 47819 **RATIONALE**  
47820 None.
- 47821 **FUTURE DIRECTIONS**  
47822 None.
- 47823 **SEE ALSO**  
47824 *open()*, the Base Definitions volume of IEEE Std 1003.1-200x, <**unistd.h**>
- 47825 **CHANGE HISTORY**  
47826 First released in Issue 4, Version 2.
- 47827 **Issue 5**  
47828 Moved from X/OPEN UNIX extension to BASE.  
47829 Large File Summit extensions are added.
- 47830 **Issue 6**  
47831 This reference page is split out from the *truncate()* reference page.  
47832 The DESCRIPTION is updated to avoid use of the term “must” for application requirements. |  
47833 The wording of the mandatory [ELOOP] error condition is updated, and a second optional  
47834 [ELOOP] error condition is added.

47835 **NAME**

47836           truncf, trunc1 — round to truncated integer value

47837 **SYNOPSIS**

47838           #include &lt;math.h&gt;

47839           float truncf(float x);

47840           long double trunc1(long double x);

47841 **DESCRIPTION**47842           Refer to *trunc()*.



47843 **NAME**

47844           tsearch — search a binary search tree

47845 **SYNOPSIS**

47846 XSI       #include &lt;search.h&gt;

47847           void \*tsearch(const void \*key, void \*\*rootp,  
47848                       int (\*compar)(const void \*, const void \*));

47849

47850 **DESCRIPTION**47851           Refer to *tdelete()*.

47852 **NAME**

47853 `ttyname`, `ttyname_r` — find pathname of a terminal

47854 **SYNOPSIS**

47855 `#include <unistd.h>`

47856 `char *ttyname(int fd);`

47857 TSF `int ttyname_r(int fd, char *name, size_t namesize);`

47858

47859 **DESCRIPTION**

47860 The `ttyname()` function shall return a pointer to a string containing a null-terminated pathname |  
 47861 of the terminal associated with file descriptor *fd*. The return value may point to static data |  
 47862 whose content is overwritten by each call.

47863 The `ttyname()` function need not be reentrant. A function that is not required to be reentrant is |  
 47864 not required to be thread-safe.

47865 TSF The `ttyname_r()` function shall store the null-terminated pathname of the terminal associated |  
 47866 with the file descriptor *fd* in the character array referenced by *name*. The array is *namesize* |  
 47867 characters long and should have space for the name and the terminating null character. The |  
 47868 maximum length of the terminal name shall be `{TTY_NAME_MAX}`.

47869 **RETURN VALUE**

47870 Upon successful completion, `ttyname()` shall return a pointer to a string. Otherwise, a null |  
 47871 pointer shall be returned and `errno` set to indicate the error.

47872 TSF If successful, the `ttyname_r()` function shall return zero. Otherwise, an error number shall be |  
 47873 returned to indicate the error.

47874 **ERRORS**

47875 The `ttyname()` function may fail if:

47876 [EBADF] The *fd* argument is not a valid file descriptor.

47877 [ENOTTY] The *fd* argument does not refer to a terminal.

47878 The `ttyname_r()` function may fail if:

47879 TSF [EBADF] The *fd* argument is not a valid file descriptor.

47880 TSF [ENOTTY] The *fd* argument does not refer to a terminal.

47881 TSF [ERANGE] The value of *namesize* is smaller than the length of the string to be returned |  
 47882 including the terminating null character.

47883 **EXAMPLES**

47884 None.

47885 **APPLICATION USAGE**

47886 None.

47887 **RATIONALE**

47888 The term *terminal* is used instead of the historical term *terminal device* in order to avoid a |  
 47889 reference to an undefined term.

47890 The thread-safe version places the terminal name in a user-supplied buffer and returns a non- |  
 47891 zero value if it fails. The non-thread-safe version may return the name in a static data area that |  
 47892 may be overwritten by each call.

47893 **FUTURE DIRECTIONS**

47894 None.

47895 **SEE ALSO**

47896 The Base Definitions volume of IEEE Std 1003.1-200x, &lt;unistd.h&gt;

47897 **CHANGE HISTORY**

47898 First released in Issue 1. Derived from Issue 1 of the SVID.

47899 **Issue 5**47900 The *ttyname\_r()* function is included for alignment with the POSIX Threads Extension.47901 A note indicating that the *ttyname()* function need not be reentrant is added to the  
47902 DESCRIPTION.47903 **Issue 6**47904 The *ttyname\_r()* function is marked as part of the Thread-Safe Functions option.47905 The following new requirements on POSIX implementations derive from alignment with the  
47906 Single UNIX Specification:

- 47907
- The statement that *errno* is set on error is added.
- 47908
- The [EBADF] and [ENOTTY] optional error conditions are added.

47909 **NAME**

47910 twalk — traverse a binary search tree

47911 **SYNOPSIS**

```
47912 xSI #include <search.h>
```

```
47913 void twalk(const void *root,  
47914            void (*action)(const void *, VISIT, int ));  
47915
```

47916 **DESCRIPTION**

47917 Refer to *tdelete()*.

47918 **NAME**

47919           tzname — timezone strings

47920 **SYNOPSIS**

47921 cx       #include &lt;time.h&gt;

47922           extern char \*tzname[2];

47923

47924 **DESCRIPTION**47925           Refer to *tzset()*.

47926 **NAME**

47927 daylight, timezone, tzname, tzset — set timezone conversion information

47928 **SYNOPSIS**

47929 #include &lt;time.h&gt;

47930 XSI extern int daylight;

47931 extern long timezone;

47932 CX extern char \*tzname[2];

47933 void tzset(void);

47934

47935 **DESCRIPTION**

47936 The *tzset()* function shall use the value of the environment variable *TZ* to set time conversion  
 47937 information used by *ctime()*, *localtime()*, *mktime()*, and *strftime()*. If *TZ* is absent from the  
 47938 environment, implementation-defined default timezone information shall be used.

47939 The *tzset()* function shall set the external variable *tzname* as follows:

47940 tzname[0] = "std";

47941 tzname[1] = "dst";

47942 where *std* and *dst* are as described in the Base Definitions volume of IEEE Std 1003.1-200x,  
 47943 Chapter 8, Environment Variables.

47944 XSI The *tzset()* function also shall set the external variable *daylight* to 0 if Daylight Savings Time  
 47945 conversions should never be applied for the timezone in use; otherwise, non-zero. The external  
 47946 variable *timezone* shall be set to the difference, in seconds, between Coordinated Universal Time  
 47947 (UTC) and local standard time.

47948 **RETURN VALUE**47949 The *tzset()* function shall not return a value.47950 **ERRORS**

47951 No errors are defined.

47952 **EXAMPLES**

47953 Example TZ variables and their timezone differences are given in the table below:

47954

	<b>TZ</b>	<i>timezone</i>
47956	EST5EDT	5*60*60
47957	GMT0	0*60*60
47958	JST-9	-9*60*60
47959	MET-1MEST	-1*60*60
47960	MST7MDT	7*60*60
47961	PST8PDT	8*60*60

47962 **APPLICATION USAGE**

47963 None.

47964 **RATIONALE**

47965 None.

47966 **FUTURE DIRECTIONS**

47967 None.

47968 **SEE ALSO**

47969 *ctime()*, *localtime()*, *mktime()*, *strftime()*, the Base Definitions volume of IEEE Std 1003.1-200x,  
47970 **<time.h>**

47971 **CHANGE HISTORY**

47972 First released in Issue 1. Derived from Issue 1 of the SVID. |

47973 **Issue 6** |

47974 The example is corrected. |

47975 **NAME**

47976           ualarm — set the interval timer

47977 **SYNOPSIS**

47978 OB XSI   #include &lt;unistd.h&gt;

47979           useconds\_t ualarm(useconds\_t *useconds*, useconds\_t *interval*);

47980

47981 **DESCRIPTION**

47982       The *ualarm()* function shall cause the SIGALRM signal to be generated for the calling process  
 47983       after the number of realtime microseconds specified by the *useconds* argument has elapsed.  
 47984       When the *interval* argument is non-zero, repeated timeout notification occurs with a period in  
 47985       microseconds specified by the *interval* argument. If the notification signal, SIGALRM, is not  
 47986       caught or ignored, the calling process is terminated.

47987       Implementations may place limitations on the granularity of timer values. For each interval  
 47988       timer, if the requested timer value requires a finer granularity than the implementation supports,  
 47989       the actual timer value shall be rounded up to the next supported value.

47990       Interactions between *ualarm()* and any of the following are unspecified:

47991            *alarm()*  
 47992            *nanosleep()*  
 47993            *setitimer()*  
 47994            *timer\_create()*  
 47995            *timer\_delete()*  
 47996            *timer\_getoverrun()*  
 47997            *timer\_gettime()*  
 47998            *timer\_settime()*  
 47999            *sleep()*

48000 **RETURN VALUE**

48001       The *ualarm()* function shall return the number of microseconds remaining from the previous  
 48002       *ualarm()* call. If no timeouts are pending or if *ualarm()* has not previously been called, *ualarm()*  
 48003       shall return 0.

48004 **ERRORS**

48005       No errors are defined.

48006 **EXAMPLES**

48007       None.

48008 **APPLICATION USAGE**

48009       Applications are recommended to use *nanosleep()* if the Timers option is supported, or  
 48010       *setitimer()*, *timer\_create()*, *timer\_delete()*, *timer\_getoverrun()*, *timer\_gettime()*, or *timer\_settime()*  
 48011       instead of this function.

48012 **RATIONALE**

48013       None.

48014 **FUTURE DIRECTIONS**

48015       None.

48016 **SEE ALSO**

48017       *alarm()*, *nanosleep()*, *setitimer()*, *sleep()*, *timer\_create()*, *timer\_delete()*, *timer\_getoverrun()*, the Base  
 48018       Definitions volume of IEEE Std 1003.1-200x, <unistd.h>



48019 **CHANGE HISTORY**

48020 First released in Issue 4, Version 2.

48021 **Issue 5**

48022 Moved from X/OPEN UNIX extension to BASE.

48023 **Issue 6**

48024 This function is marked obsolescent.

48025 **NAME**

48026 ulimit — get and set process limits

48027 **SYNOPSIS**48028 XSI `#include <ulimit.h>`48029 `long ulimit(int cmd, ...);`

48030

48031 **DESCRIPTION**

48032 The *ulimit()* function shall control process limits. The process limits that can be controlled by  
 48033 this function include the maximum size of a single file that can be written (this is equivalent to  
 48034 using *setrlimit()* with `RLIMIT_FSIZE`). The *cmd* values, defined in `<ulimit.h>` include:

48035 `UL_GETFSIZE` Return the file size limit (`RLIMIT_FSIZE`) of the process. The limit shall be in  
 48036 units of 512-byte blocks and shall be inherited by child processes. Files of any  
 48037 size can be read. The return value shall be the integer part of the soft file size  
 48038 limit divided by 512. If the result cannot be represented as a **long**, the result is  
 48039 unspecified.

48040 `UL_SETFSIZE` Set the file size limit for output operations of the process to the value of the  
 48041 second argument, taken as a **long**, multiplied by 512. If the result would  
 48042 overflow an `rlim_t`, the actual value set is unspecified. Any process may  
 48043 decrease its own limit, but only a process with appropriate privileges may  
 48044 increase the limit. The return value shall be the integer part of the new file size  
 48045 limit divided by 512.

48046 The *ulimit()* function shall not change the setting of *errno* if successful.

48047 As all return values are permissible in a successful situation, an application wishing to check for  
 48048 error situations should set *errno* to 0, then call *ulimit()*, and, if it returns `-1`, check to see if *errno* is  
 48049 non-zero.

48050 **RETURN VALUE**

48051 Upon successful completion, *ulimit()* shall return the value of the requested limit. Otherwise, `-1`  
 48052 shall be returned and *errno* set to indicate the error.

48053 **ERRORS**

48054 The *ulimit()* function shall fail and the limit shall be unchanged if:

48055 `[EINVAL]` The *cmd* argument is not valid.

48056 `[EPERM]` A process not having appropriate privileges attempts to increase its file size  
 48057 limit.

48058 **EXAMPLES**

48059 None.

48060 **APPLICATION USAGE**

48061 None.

48062 **RATIONALE**

48063 None.

48064 **FUTURE DIRECTIONS**

48065 None.

48066 **SEE ALSO**

48067 *getrlimit()*, *setrlimit()*, *write()*, the Base Definitions volume of IEEE Std 1003.1-200x, <ulimit.h>

48068 **CHANGE HISTORY**

48069 First released in Issue 1. Derived from Issue 1 of the SVID.

48070 **Issue 5**

48071 In the description of UL\_SETFSIZE, the text is corrected to refer to **rlim\_t** rather than the spurious **rlimit\_t**.

48073 The DESCRIPTION is updated to indicate that *errno* is not changed if the function is successful.

48074 **NAME**

48075 umask — set and get file mode creation mask

48076 **SYNOPSIS**

48077 #include <sys/stat.h>

48078 mode\_t umask(mode\_t *cmask*);

48079 **DESCRIPTION**

48080 The *umask()* function shall set the process' file mode creation mask to *cmask* and return the  
48081 previous value of the mask. Only the file permission bits of *cmask* (see <sys/stat.h>) are used; the  
48082 meaning of the other bits is implementation-defined.

48083 The process' file mode creation mask is used during *open()*, *creat()*, *mkdir()*, and *mkfifo()* to turn  
48084 off permission bits in the *mode* argument supplied. Bit positions that are set in *cmask* are cleared  
48085 in the mode of the created file.

48086 **RETURN VALUE**

48087 The file permission bits in the value returned by *umask()* shall be the previous value of the file  
48088 mode creation mask. The state of any other bits in that value is unspecified, except that a  
48089 subsequent call to *umask()* with the returned value as *cmask* shall leave the state of the mask the  
48090 same as its state before the first call, including any unspecified use of those bits.

48091 **ERRORS**

48092 No errors are defined.

48093 **EXAMPLES**

48094 None.

48095 **APPLICATION USAGE**

48096 None.

48097 **RATIONALE**

48098 Unsigned argument and return types for *umask()* were proposed. The return type and the  
48099 argument were both changed to **mode\_t**.

48100 Historical implementations have made use of additional bits in *cmask* for their implementation-  
48101 defined purposes. The addition of the text that the meaning of other bits of the field is  
48102 implementation-defined permits these implementations to conform to this volume of  
48103 IEEE Std 1003.1-200x.

48104 **FUTURE DIRECTIONS**

48105 None.

48106 **SEE ALSO**

48107 *creat()*, *mkdir()*, *mkfifo()*, *open()*, the Base Definitions volume of IEEE Std 1003.1-200x,  
48108 <sys/stat.h>, <sys/types.h>

48109 **CHANGE HISTORY**

48110 First released in Issue 1. Derived from Issue 1 of the SVID.

48111 **Issue 6**

48112 In the SYNOPSIS, the optional include of the <sys/types.h> header is removed. |

- 48113 The following new requirements on POSIX implementations derive from alignment with the |  
48114 Single UNIX Specification:
- 48115 • The requirement to include `<sys/types.h>` has been removed. Although `<sys/types.h>` was  
48116 required for conforming implementations of previous POSIX specifications, it was not  
48117 required for UNIX applications.

48118 **NAME**

48119            **uname** — get name of current system

48120 **SYNOPSIS**

48121            #include <sys/utsname.h>

48122            int uname(struct utsname \*name);

48123 **DESCRIPTION**

48124            The *uname()* function shall store information identifying the current system in the structure pointed to by *name*.

48126            The *uname()* function uses the **utsname** structure defined in <sys/utsname.h>.

48127            The *uname()* function shall return a string naming the current system in the character array *sysname*. Similarly, *nodename* shall contain the name of this node within an implementation-defined communications network. The arrays *release* and *version* shall further identify the operating system. The array *machine* shall contain a name that identifies the hardware that the system is running on.

48132            The format of each member is implementation-defined.

48133 **RETURN VALUE**

48134            Upon successful completion, a non-negative value shall be returned. Otherwise, -1 shall be returned and *errno* set to indicate the error.

48136 **ERRORS**

48137            No errors are defined.

48138 **EXAMPLES**

48139            None.

48140 **APPLICATION USAGE**

48141            The inclusion of the *nodename* member in this structure does not imply that it is sufficient information for interfacing to communications networks.

48143 **RATIONALE**

48144            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>.

48147            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.

48151            The *uname()* function originated in 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.

48154            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 former functions are included in this specification, the latter are not.

48158 **FUTURE DIRECTIONS**

48159            None.

48160 **SEE ALSO**

48161           The Base Definitions volume of IEEE Std 1003.1-200x, <sys/utsname.h>

48162 **CHANGE HISTORY**

48163           First released in Issue 1. Derived from Issue 1 of the SVID.

48164 **NAME**

48165 ungetc — push byte back into input stream

48166 **SYNOPSIS**

48167 #include &lt;stdio.h&gt;

48168 int ungetc(int *c*, FILE \**stream*);48169 **DESCRIPTION**

48170 cx The functionality described on this reference page is aligned with the ISO C standard. Any  
48171 conflict between the requirements described here and the ISO C standard is unintentional. This  
48172 volume of IEEE Std 1003.1-200x defers to the ISO C standard.

48173 The *ungetc()* function shall push the byte specified by *c* (converted to an **unsigned char**) back  
48174 onto the input stream pointed to by *stream*. The pushed-back bytes shall be returned by  
48175 subsequent reads on that stream in the reverse order of their pushing. A successful intervening  
48176 call (with the stream pointed to by *stream*) to a file-positioning function (*fseek()*, *fsetpos()*, or  
48177 *rewind()*) shall discard any pushed-back bytes for the stream. The external storage  
48178 corresponding to the stream shall be unchanged.

48179 One byte of push-back shall be provided. If *ungetc()* is called too many times on the same stream  
48180 without an intervening read or file-positioning operation on that stream, the operation may fail.

48181 If the value of *c* equals that of the macro EOF, the operation shall fail and the input stream shall  
48182 be left unchanged.

48183 A successful call to *ungetc()* shall clear the end-of-file indicator for the stream. The value of the  
48184 file-position indicator for the stream after reading or discarding all pushed-back bytes shall be  
48185 the same as it was before the bytes were pushed back. The file-position indicator is decremented  
48186 by each successful call to *ungetc()*; if its value was 0 before a call, its value is unspecified after  
48187 the call.

48188 **RETURN VALUE**

48189 Upon successful completion, *ungetc()* shall return the byte pushed back after conversion.  
48190 Otherwise, it shall return EOF.

48191 **ERRORS**

48192 No errors are defined.

48193 **EXAMPLES**

48194 None.

48195 **APPLICATION USAGE**

48196 None.

48197 **RATIONALE**

48198 None.

48199 **FUTURE DIRECTIONS**

48200 None.

48201 **SEE ALSO**

48202 *fseek()*, *getc()*, *fsetpos()*, *read()*, *rewind()*, *setbuf()*, the Base Definitions volume of  
48203 IEEE Std 1003.1-200x, <stdio.h>

48204 **CHANGE HISTORY**

48205 First released in Issue 1. Derived from Issue 1 of the SVID.



48206 **NAME**

48207 ungetwc — push wide-character code back into input stream

48208 **SYNOPSIS**

48209 #include &lt;stdio.h&gt;

48210 #include &lt;wchar.h&gt;

48211 wint\_t ungetwc(wint\_t *wc*, FILE \**stream*);48212 **DESCRIPTION**

48213 **CX** The functionality described on this reference page is aligned with the ISO C standard. Any  
 48214 conflict between the requirements described here and the ISO C standard is unintentional. This  
 48215 volume of IEEE Std 1003.1-200x defers to the ISO C standard.

48216 The *ungetwc()* function shall push the character corresponding to the wide-character code  
 48217 specified by *wc* back onto the input stream pointed to by *stream*. The pushed-back characters  
 48218 shall be returned by subsequent reads on that stream in the reverse order of their pushing. A  
 48219 successful intervening call (with the stream pointed to by *stream*) to a file-positioning function  
 48220 (*fseek()*, *fsetpos()*, or *rewind()*) discards any pushed-back characters for the stream. The external  
 48221 storage corresponding to the stream is unchanged.

48222 At least one character of push-back shall be provided. If *ungetwc()* is called too many times on  
 48223 the same stream without an intervening read or file-positioning operation on that stream, the  
 48224 operation may fail.

48225 If the value of *wc* equals that of the macro WEOF, the operation shall fail and the input stream  
 48226 shall be left unchanged.

48227 A successful call to *ungetwc()* shall clear the end-of-file indicator for the stream. The value of the  
 48228 file-position indicator for the stream after reading or discarding all pushed-back characters shall  
 48229 be the same as it was before the characters were pushed back. The file-position indicator is  
 48230 decremented (by one or more) by each successful call to *ungetwc()*; if its value was 0 before a  
 48231 call, its value is unspecified after the call.

48232 **RETURN VALUE**

48233 Upon successful completion, *ungetwc()* shall return the wide-character code corresponding to  
 48234 the pushed-back character. Otherwise, it shall return WEOF.

48235 **ERRORS**48236 The *ungetwc()* function may fail if:

48237 **CX** [EILSEQ] An invalid character sequence is detected, or a wide-character code does not  
 48238 correspond to a valid character.

48239 **EXAMPLES**

48240 None.

48241 **APPLICATION USAGE**

48242 None.

48243 **RATIONALE**

48244 None.

48245 **FUTURE DIRECTIONS**

48246 None.

48247 **SEE ALSO**

48248 *fseek()*, *fsetpos()*, *read()*, *rewind()*, *setbuf()*, the Base Definitions volume of IEEE Std 1003.1-200x,  
48249 `<stdio.h>`, `<wchar.h>`

48250 **CHANGE HISTORY**

48251 First released in Issue 4. Derived from the MSE working draft.

48252 **Issue 5**

48253 The Optional Header (OH) marking is removed from `<stdio.h>`. |

48254 **Issue 6**

48255 The [EILSEQ] optional error condition is marked CX. |

48256 **NAME**

48257 unlink — remove a directory entry

48258 **SYNOPSIS**

48259 #include &lt;unistd.h&gt;

48260 int unlink(const char \*path);

48261 **DESCRIPTION**

48262 The *unlink()* function shall remove a link to a file. If *path* names a symbolic link, *unlink()* shall  
 48263 remove the symbolic link named by *path* and shall not affect any file or directory named by the  
 48264 contents of the symbolic link. Otherwise, *unlink()* shall remove the link named by the pathname  
 48265 pointed to by *path* and shall decrement the link count of the file referenced by the link.

48266 When the file's link count becomes 0 and no process has the file open, the space occupied by the  
 48267 file shall be freed and the file shall no longer be accessible. If one or more processes have the file  
 48268 open when the last link is removed, the link shall be removed before *unlink()* returns, but the  
 48269 removal of the file contents shall be postponed until all references to the file are closed.

48270 The *path* argument shall not name a directory unless the process has appropriate privileges and  
 48271 the implementation supports using *unlink()* on directories.

48272 Upon successful completion, *unlink()* shall mark for update the *st\_ctime* and *st\_mtime* fields of  
 48273 the parent directory. Also, if the file's link count is not 0, the *st\_ctime* field of the file shall be  
 48274 marked for update.

48275 **RETURN VALUE**

48276 Upon successful completion, 0 shall be returned. Otherwise, -1 shall be returned and *errno* set to  
 48277 indicate the error. If -1 is returned, the named file shall not be changed.

48278 **ERRORS**48279 The *unlink()* function shall fail and shall not unlink the file if:

48280 [EACCES] Search permission is denied for a component of the path prefix, or write  
 48281 permission is denied on the directory containing the directory entry to be  
 48282 removed.

48283 [EBUSY] The file named by the *path* argument cannot be unlinked because it is being  
 48284 used by the system or another process and the implementation considers this  
 48285 an error.

48286 [ELOOP] A loop exists in symbolic links encountered during resolution of the *path*  
 48287 argument.

48288 [ENAMETOOLONG] The length of the *path* argument exceeds {PATH\_MAX} or a pathname  
 48289 component is longer than {NAME\_MAX}.  
 48290

48291 [ENOENT] A component of *path* does not name an existing file or *path* is an empty string.

48292 [ENOTDIR] A component of the path prefix is not a directory.

48293 [EPERM] The file named by *path* is a directory, and either the calling process does not  
 48294 have appropriate privileges, or the implementation prohibits using *unlink()*  
 48295 on directories.

48296 XSI [EPERM] or [EACCES]

48297 The S\_ISVTX flag is set on the directory containing the file referred to by the  
 48298 *path* argument and the caller is not the file owner, nor is the caller the  
 48299 directory owner, nor does the caller have appropriate privileges.

48300 [EROFS] The directory entry to be unlinked is part of a read-only file system.

48301 The *unlink()* function may fail and not unlink the file if:

48302 XSI [EBUSY] The file named by *path* is a named STREAM.

48303 [ELOOP] More than {SYMLOOP\_MAX} symbolic links were encountered during  
48304 resolution of the *path* argument.

48305 [ENAMETOOLONG]  
48306 As a result of encountering a symbolic link in resolution of the *path* argument, |  
48307 the length of the substituted pathname string exceeded {PATH\_MAX}. |

48308 [ETXTBSY] The entry to be unlinked is the last directory entry to a pure procedure (shared  
48309 text) file that is being executed.

48310 **EXAMPLES**

48311 **Removing a Link to a File**

48312 The following example shows how to remove a link to a file named `/home/cnd/mod1` by  
48313 removing the entry named `/modules/pass1`.

```
48314 #include <unistd.h>
48315 char *path = "/modules/pass1";
48316 int status;
48317 ...
48318 status = unlink(path);
```

48319 **Checking for an Error**

48320 The following example fragment creates a temporary password lock file named **LOCKFILE**,  
48321 which is defined as `/etc/ptmp`, and gets a file descriptor for it. If the file cannot be opened for  
48322 writing, *unlink()* is used to remove the link between the file descriptor and **LOCKFILE**.

```
48323 #include <sys/types.h>
48324 #include <stdio.h>
48325 #include <fcntl.h>
48326 #include <errno.h>
48327 #include <unistd.h>
48328 #include <sys/stat.h>
48329 #define LOCKFILE "/etc/ptmp"
48330 int pfd; /* Integer for file descriptor returned by open call. */
48331 FILE *fpfd; /* File pointer for use in putpwent(). */
48332 ...
48333 /* Open password Lock file. If it exists, this is an error. */
48334 if ((pfd = open(LOCKFILE, O_WRONLY | O_CREAT | O_EXCL, S_IRUSR
48335 | S_IWUSR | S_IRGRP | S_IROTH)) == -1) {
48336     fprintf(stderr, "Cannot open /etc/ptmp. Try again later.\n");
48337     exit(1);
48338 }
48339 /* Lock file created, proceed with fdopen of lock file so that
48340 putpwent() can be used.
48341 */
48342 if ((fpfd = fdopen(pfd, "w")) == NULL) {
```

```

48343         close(pfd);
48344         unlink(LOCKFILE);
48345         exit(1);
48346     }

```

### 48347 **Replacing Files**

48348 The following example fragment uses *unlink()* to discard links to files, so that they can be  
48349 replaced with new versions of the files. The first call remove the link to **LOCKFILE** if an error  
48350 occurs. Successive calls remove the links to **SAVEFILE** and **PASSWDFILE** so that new links can  
48351 be created, then removes the link to **LOCKFILE** when it is no longer needed.

```

48352     #include <sys/types.h>
48353     #include <stdio.h>
48354     #include <fcntl.h>
48355     #include <errno.h>
48356     #include <unistd.h>
48357     #include <sys/stat.h>

48358     #define LOCKFILE "/etc/ptmp"
48359     #define PASSWDFILE "/etc/passwd"
48360     #define SAVEFILE "/etc/opasswd"
48361     ...
48362     /* If no change was made, assume error and leave passwd unchanged. */
48363     if (!valid_change) {
48364         fprintf(stderr, "Could not change password for user %s\n", user);
48365         unlink(LOCKFILE);
48366         exit(1);
48367     }

48368     /* Change permissions on new password file. */
48369     chmod(LOCKFILE, S_IRUSR | S_IRGRP | S_IROTH);

48370     /* Remove saved password file. */
48371     unlink(SAVEFILE);

48372     /* Save current password file. */
48373     link(PASSWDFILE, SAVEFILE);

48374     /* Remove current password file. */
48375     unlink(PASSWDFILE);

48376     /* Save new password file as current password file. */
48377     link(LOCKFILE, PASSWDFILE);

48378     /* Remove lock file. */
48379     unlink(LOCKFILE);

48380     exit(0);

```

### 48381 **APPLICATION USAGE**

48382 Applications should use *rmdir()* to remove a directory.

### 48383 **RATIONALE**

48384 Unlinking a directory is restricted to the superuser in many historical implementations for  
48385 reasons given in *link()* (see also *rename()*).

48386 The meaning of [EBUSY] in historical implementations is “mount point busy”. Since this volume  
 48387 of IEEE Std 1003.1-200x does not cover the system administration concepts of mounting and  
 48388 unmounting, the description of the error was changed to “resource busy”. (This meaning is used  
 48389 by some device drivers when a second process tries to open an exclusive use device.) The  
 48390 wording is also intended to allow implementations to refuse to remove a directory if it is the  
 48391 root or current working directory of any process.

48392 **FUTURE DIRECTIONS**

48393 None.

48394 **SEE ALSO**

48395 *close()*, *link()*, *remove()*, *rmdir()*, the Base Definitions volume of IEEE Std 1003.1-200x,  
 48396 <unistd.h>

48397 **CHANGE HISTORY**

48398 First released in Issue 1. Derived from Issue 1 of the SVID.

48399 **Issue 5**

48400 The [EBUSY] error is added to the “may fail” part of the ERRORS section.

48401 **Issue 6**

48402 The following new requirements on POSIX implementations derive from alignment with the  
 48403 Single UNIX Specification:

- 48404 • In the DESCRIPTION, the effect is specified if *path* specifies a symbolic link.
- 48405 • The [ELOOP] mandatory error condition is added.
- 48406 • A second [ENAMETOOLONG] is added as an optional error condition.
- 48407 • The [ETXTBSY] optional error condition is added.

48408 The following changes were made to align with the IEEE P1003.1a draft standard:

- 48409 • The [ELOOP] optional error condition is added.

48410 The DESCRIPTION is updated to avoid use of the term “must” for application requirements.

48411 **NAME**

48412 unlockpt — unlock a pseudo-terminal master/slave pair

48413 **SYNOPSIS**48414 XSI `#include <stdlib.h>`48415 `int unlockpt(int fildev);`

48416

48417 **DESCRIPTION**48418 The `unlockpt()` function shall unlock the slave pseudo-terminal device associated with the  
48419 master to which *fildev* refers.48420 Conforming applications shall ensure that they call `unlockpt()` before opening the slave side of a  
48421 pseudo-terminal device.48422 **RETURN VALUE**48423 Upon successful completion, `unlockpt()` shall return 0. Otherwise, it shall return `-1` and set *errno*  
48424 to indicate the error.48425 **ERRORS**48426 The `unlockpt()` function may fail if:48427 [EBADF] The *fildev* argument is not a file descriptor open for writing.48428 [EINVAL] The *fildev* argument is not associated with a master pseudo-terminal device.48429 **EXAMPLES**

48430 None.

48431 **APPLICATION USAGE**

48432 None.

48433 **RATIONALE**

48434 None.

48435 **FUTURE DIRECTIONS**

48436 None.

48437 **SEE ALSO**48438 `grantpt()`, `open()`, `ptsname()`, the Base Definitions volume of IEEE Std 1003.1-200x, `<stdlib.h>`48439 **CHANGE HISTORY**

48440 First released in Issue 4, Version 2.

48441 **Issue 5**

48442 Moved from X/OPEN UNIX extension to BASE.

48443 **Issue 6**

48444 The DESCRIPTION is updated to avoid use of the term “must” for application requirements.

48445 **NAME**

48446           unsetenv — remove environment variable

48447 **SYNOPSIS**48448 `cx`       #include <stdlib.h>

48449           int unsetenv(const char \*name);

48450

48451 **DESCRIPTION**

48452       The *unsetenv()* function shall remove an environment variable from the environment of the  
48453       calling process. The *name* argument points to a string, which is the name of the variable to be  
48454       removed. The named argument shall not contain an '=' character. If the named variable does  
48455       not exist in the current environment, the environment shall be unchanged and the function is  
48456       considered to have completed successfully.

48457       If the application modifies *environ* or the pointers to which it points, the behavior of *unsetenv()* is  
48458       undefined. The *unsetenv()* function shall update the list of pointers to which *environ* points.

48459       The *unsetenv()* function need not be reentrant. A function that is not required to be reentrant is  
48460       not required to be thread-safe.

48461 **RETURN VALUE**

48462       Upon successful completion, zero shall be returned. Otherwise, -1 shall be returned, *errno* set to  
48463       indicate the error, and the environment shall be unchanged.

48464 **ERRORS**48465       The *unsetenv()* function shall fail if:

48466       [EINVAL]       The *name* argument is a null pointer, points to an empty string, or points to a  
48467       string containing an '=' character.

48468 **EXAMPLES**

48469       None.

48470 **APPLICATION USAGE**

48471       None.

48472 **RATIONALE**48473       Refer to the RATIONALE section in *setenv()*.48474 **FUTURE DIRECTIONS**

48475       None.

48476 **SEE ALSO**

48477       *getenv()*, *setenv()*, the Base Definitions volume of IEEE Std 1003.1-200x, <stdlib.h>,  
48478       <sys/types.h>, <unistd.h>

48479 **CHANGE HISTORY**

48480       First released in Issue 6. Derived from the IEEE P1003.1a draft standard.



48481 **NAME**

48482            usleep — suspend execution for an interval

48483 **SYNOPSIS**

48484 OB XSI    #include &lt;unistd.h&gt;

48485            int usleep(useconds\_t useconds);

48486

48487 **DESCRIPTION**

48488        The *usleep()* function shall cause the calling thread to be suspended from execution until either  
 48489        the number of realtime microseconds specified by the argument *useconds* has elapsed or a signal  
 48490        is delivered to the calling thread and its action is to invoke a signal-catching function or to  
 48491        terminate the process. The suspension time may be longer than requested due to the scheduling  
 48492        of other activity by the system.

48493        The *useconds* argument shall be less than one million. If the value of *useconds* is 0, then the call  
 48494        has no effect.

48495        If a SIGALRM signal is generated for the calling process during execution of *usleep()* and if the  
 48496        SIGALRM signal is being ignored or blocked from delivery, it is unspecified whether *usleep()*  
 48497        returns when the SIGALRM signal is scheduled. If the signal is being blocked, it is also  
 48498        unspecified whether it remains pending after *usleep()* returns or it is discarded.

48499        If a SIGALRM signal is generated for the calling process during execution of *usleep()*, except as a  
 48500        result of a prior call to *alarm()*, and if the SIGALRM signal is not being ignored or blocked from  
 48501        delivery, it is unspecified whether that signal has any effect other than causing *usleep()* to return.

48502        If a signal-catching function interrupts *usleep()* and examines or changes either the time a  
 48503        SIGALRM is scheduled to be generated, the action associated with the SIGALRM signal, or  
 48504        whether the SIGALRM signal is blocked from delivery, the results are unspecified.

48505        If a signal-catching function interrupts *usleep()* and calls *siglongjmp()* or *longjmp()* to restore an  
 48506        environment saved prior to the *usleep()* call, the action associated with the SIGALRM signal and  
 48507        the time at which a SIGALRM signal is scheduled to be generated are unspecified. It is also  
 48508        unspecified whether the SIGALRM signal is blocked, unless the process' signal mask is restored  
 48509        as part of the environment.

48510        Implementations may place limitations on the granularity of timer values. For each interval  
 48511        timer, if the requested timer value requires a finer granularity than the implementation supports,  
 48512        the actual timer value shall be rounded up to the next supported value.

48513        Interactions between *usleep()* and any of the following are unspecified:

48514            *nanosleep()*48515            *setitimer()*48516            *timer\_create()*48517            *timer\_delete()*48518            *timer\_getoverrun()*48519            *timer\_gettime()*48520            *timer\_settime()*48521            *ualarm()*48522            *sleep()*

## 48523 RETURN VALUE

48524       Upon successful completion, *usleep()* shall return 0; otherwise, it shall return -1 and set *errno* to  
48525       indicate the error.

## 48526 ERRORS

48527       The *usleep()* function may fail if:

48528       [EINVAL]       The time interval specified one million or more microseconds.

## 48529 EXAMPLES

48530       None.

## 48531 APPLICATION USAGE

48532       Applications are recommended to use *nanosleep()* if the Timers option is supported, or  
48533       *setitimer()*, *timer\_create()*, *timer\_delete()*, *timer\_getoverrun()*, *timer\_gettime()*, or *timer\_settime()*  
48534       instead of this function.

## 48535 RATIONALE

48536       None.

## 48537 FUTURE DIRECTIONS

48538       None.

## 48539 SEE ALSO

48540       *alarm()*, *getitimer()*, *nanosleep()*, *sigaction()*, *sleep()*, *timer\_create()*, *timer\_delete()*,  
48541       *timer\_getoverrun()*, the Base Definitions volume of IEEE Std 1003.1-200x, <**unistd.h**>

## 48542 CHANGE HISTORY

48543       First released in Issue 4, Version 2.

### 48544 Issue 5

48545       Moved from X/OPEN UNIX extension to BASE.

48546       The DESCRIPTION is changed to indicate that timers are now thread-based rather than  
48547       process-based.

### 48548 Issue 6

48549       The DESCRIPTION is updated to avoid use of the term “must” for application requirements.

48550       This function is marked obsolescent.

48551 **NAME**

48552 utime — set file access and modification times

48553 **SYNOPSIS**

48554 #include &lt;utime.h&gt;

48555 int utime(const char \*path, const struct utimbuf \*times);

48556 **DESCRIPTION**48557 The *utime()* function shall set the access and modification times of the file named by the *path*  
48558 argument.48559 If *times* is a null pointer, the access and modification times of the file shall be set to the current |  
48560 time. The effective user ID of the process shall match the owner of the file, or the process has |  
48561 write permission to the file or has appropriate privileges, to use *utime()* in this manner. |48562 If *times* is not a null pointer, *times* shall be interpreted as a pointer to a **utimbuf** structure and the |  
48563 access and modification times shall be set to the values contained in the designated structure. |  
48564 Only a process with effective user ID equal to the user ID of the file or a process with |  
48565 appropriate privileges may use *utime()* this way. |48566 The **utimbuf** structure is defined in the <**utime.h**> header. The times in the structure **utimbuf** |  
48567 are measured in seconds since the Epoch.48568 Upon successful completion, *utime()* shall mark the time of the last file status change, *st\_ctime*,  
48569 to be updated; see <**sys/stat.h**>.48570 **RETURN VALUE**48571 Upon successful completion, 0 shall be returned. Otherwise, -1 shall be returned and *errno* shall  
48572 be set to indicate the error, and the file times shall not be affected.48573 **ERRORS**48574 The *utime()* function shall fail if:48575 [EACCES] Search permission is denied by a component of the path prefix; or the *times*  
48576 argument is a null pointer and the effective user ID of the process does not  
48577 match the owner of the file, the process does not have write permission for the  
48578 file, and the process does not have appropriate privileges.48579 [ELOOP] A loop exists in symbolic links encountered during resolution of the *path*  
48580 argument.48581 [ENAMETOOLONG]  
48582 The length of the *path* argument exceeds {PATH\_MAX} or a pathname |  
48583 component is longer than {NAME\_MAX}. |48584 [ENOENT] A component of *path* does not name an existing file or *path* is an empty string.

48585 [ENOTDIR] A component of the path prefix is not a directory.

48586 [EPERM] The *times* argument is not a null pointer and the calling process' effective user  
48587 ID does not match the owner of the file and the calling process does not have  
48588 the appropriate privileges.

48589 [EROFS] The file system containing the file is read-only.

48590 The *utime()* function may fail if:48591 [ELOOP] More than {SYMLOOP\_MAX} symbolic links were encountered during  
48592 resolution of the *path* argument.

48593 [ENAMETOOLONG]  
48594 As a result of encountering a symbolic link in resolution of the *path* argument, |  
48595 the length of the substituted pathname string exceeded {PATH\_MAX}. |

**48596 EXAMPLES**

48597 None.

**48598 APPLICATION USAGE**

48599 None.

**48600 RATIONALE**

48601 The *actime* structure member must be present so that an application may set it, even though an  
48602 implementation may ignore it and not change the access time on the file. If an application  
48603 intends to leave one of the times of a file unchanged while changing the other, it should use  
48604 *stat()* to retrieve the file's *st\_atime* and *st\_mtime* parameters, set *actime* and *modtime* in the buffer,  
48605 and change one of them before making the *utime()* call.

**48606 FUTURE DIRECTIONS**

48607 None.

**48608 SEE ALSO**

48609 The Base Definitions volume of IEEE Std 1003.1-200x, <**sys/types.h**>, <**utime.h**>

**48610 CHANGE HISTORY**

48611 First released in Issue 1. Derived from Issue 1 of the SVID.

**48612 Issue 6**

48613 In the SYNOPSIS, the optional include of the <**sys/types.h**> header is removed.

48614 The following new requirements on POSIX implementations derive from alignment with the |  
48615 Single UNIX Specification:

48616 • The requirement to include <**sys/types.h**> has been removed. Although <**sys/types.h**> was  
48617 required for conforming implementations of previous POSIX specifications, it was not  
48618 required for UNIX applications.

48619 • The [ELOOP] mandatory error condition is added.

48620 • A second [ENAMETOOLONG] is added as an optional error condition.

48621 The following changes were made to align with the IEEE P1003.1a draft standard:

48622 • The [ELOOP] optional error condition is added.

48623 The DESCRIPTION is updated to avoid use of the term “must” for application requirements.

48624 **NAME**48625 utimes — set file access and modification times (**LEGACY**)48626 **SYNOPSIS**48627 XSI `#include <sys/time.h>`48628 `int utimes(const char *path, const struct timeval times[2]);`

48629

48630 **DESCRIPTION**

48631 The *utimes()* function shall set the access and modification times of the file pointed to by the *path*  
 48632 argument to the value of the *times* argument. The *utimes()* function allows time specifications  
 48633 accurate to the microsecond.

48634 For *utimes()*, the *times* argument is an array of **timeval** structures. The first array member  
 48635 represents the date and time of last access, and the second member represents the date and time  
 48636 of last modification. The times in the **timeval** structure are measured in seconds and  
 48637 microseconds since the Epoch, although rounding toward the nearest second may occur.

48638 If the *times* argument is a null pointer, the access and modification times of the file shall be set to |  
 48639 the current time. The effective user ID of the process shall match the owner of the file, or has |  
 48640 write access to the file or appropriate privileges to use this call in this manner. Upon completion, |  
 48641 *utimes()* shall mark the time of the last file status change, *st\_ctime*, for update.

48642 **RETURN VALUE**

48643 Upon successful completion, 0 shall be returned. Otherwise, -1 shall be returned and *errno* shall  
 48644 be set to indicate the error, and the file times shall not be affected.

48645 **ERRORS**48646 The *utimes()* function shall fail if:

48647 [EACCES] Search permission is denied by a component of the path prefix; or the *times*  
 48648 argument is a null pointer and the effective user ID of the process does not  
 48649 match the owner of the file and write access is denied.

48650 [ELOOP] A loop exists in symbolic links encountered during resolution of the *path*  
 48651 argument.

48652 [ENAMETOOLONG]  
 48653 The length of the *path* argument exceeds {PATH\_MAX} or a pathname |  
 48654 component is longer than {NAME\_MAX}.

48655 [ENOENT] A component of *path* does not name an existing file or *path* is an empty string.

48656 [ENOTDIR] A component of the path prefix is not a directory.

48657 [EPERM] The *times* argument is not a null pointer and the calling process' effective user  
 48658 ID has write access to the file but does not match the owner of the file and the  
 48659 calling process does not have the appropriate privileges.

48660 [EROFS] The file system containing the file is read-only.

48661 The *utimes()* function may fail if:

48662 [ELOOP] More than {SYMLOOP\_MAX} symbolic links were encountered during  
 48663 resolution of the *path* argument.

48664 [ENAMETOOLONG]  
 48665 Pathname resolution of a symbolic link produced an intermediate result |  
 48666 whose length exceeds {PATH\_MAX}.

48667 **EXAMPLES**

48668           None.

48669 **APPLICATION USAGE**48670           For applications portability, the *utime()* function should be used to set file access and  
48671           modification times instead of *utimes()*.48672 **RATIONALE**

48673           None.

48674 **FUTURE DIRECTIONS**

48675           This function may be withdrawn in a future version.

48676 **SEE ALSO**48677           The Base Definitions volume of IEEE Std 1003.1-200x, <**sys/time.h**>48678 **CHANGE HISTORY**

48679           First released in Issue 4, Version 2.

48680 **Issue 5**

48681           Moved from X/OPEN UNIX extension to BASE.

48682 **Issue 6**

48683           This function is marked LEGACY.

48684           The DESCRIPTION is updated to avoid use of the term “must” for application requirements. |

48685           The wording of the mandatory [ELOOP] error condition is updated, and a second optional  
48686           [ELOOP] error condition is added.

48687 **NAME**

48688       va\_arg, va\_copy, va\_end, va\_start — handle variable argument list

48689 **SYNOPSIS**

48690       #include &lt;stdarg.h&gt;

48691       type va\_arg(va\_list ap, type);

48692       void va\_copy(va\_list dest, va\_list src);

48693       void va\_end(va\_list ap);

48694       void va\_start(va\_list ap, argN);

48695 **DESCRIPTION**

48696       Refer to the Base Definitions volume of IEEE Std 1003.1-200x, &lt;stdarg.h&gt;.

48697 **NAME**

48698 vfork — create new process; share virtual memory

48699 **SYNOPSIS**

48700 OB XSI #include &lt;unistd.h&gt;

48701 pid\_t vfork(void);

48702

48703 **DESCRIPTION**

48704 The *vfork()* function shall be equivalent to *fork()*, except that the behavior is undefined if the  
 48705 process created by *vfork()* either modifies any data other than a variable of type **pid\_t** used to  
 48706 store the return value from *vfork()*, or returns from the function in which *vfork()* was called, or  
 48707 calls any other function before successfully calling *\_exit()* or one of the *exec* family of functions.

48708 **RETURN VALUE**

48709 Upon successful completion, *vfork()* shall return 0 to the child process and return the process ID  
 48710 of the child process to the parent process. Otherwise, -1 shall be returned to the parent, no child  
 48711 process shall be created, and *errno* shall be set to indicate the error.

48712 **ERRORS**48713 The *vfork()* function shall fail if:

48714 [EAGAIN] The system-wide limit on the total number of processes under execution  
 48715 would be exceeded, or the system-imposed limit on the total number of  
 48716 processes under execution by a single user would be exceeded.

48717 [ENOMEM] There is insufficient swap space for the new process.

48718 **EXAMPLES**

48719 None.

48720 **APPLICATION USAGE**

48721 Conforming applications are recommended not to depend on *vfork()*, but to use *fork()* instead.  
 48722 The *vfork()* function may be withdrawn in a future version.

48723 On some implementations, *vfork()* is equivalent to *fork()*.

48724 The *vfork()* function differs from *fork()* only in that the child process can share code and data  
 48725 with the calling process (parent process). This speeds cloning activity significantly at a risk to  
 48726 the integrity of the parent process if *vfork()* is misused.

48727 The use of *vfork()* for any purpose except as a prelude to an immediate call to a function from  
 48728 the *exec* family, or to *\_exit()*, is not advised.

48729 The *vfork()* function can be used to create new processes without fully copying the address  
 48730 space of the old process. If a forked process is simply going to call *exec*, the data space copied  
 48731 from the parent to the child by *fork()* is not used. This is particularly inefficient in a paged  
 48732 environment, making *vfork()* particularly useful. Depending upon the size of the parent's data  
 48733 space, *vfork()* can give a significant performance improvement over *fork()*.

48734 The *vfork()* function can normally be used just like *fork()*. It does not work, however, to return  
 48735 while running in the child's context from the caller of *vfork()* since the eventual return from  
 48736 *vfork()* would then return to a no longer existent stack frame. Care should be taken, also, to call  
 48737 *\_exit()* rather than *exit()* if *exec* cannot be used, since *exit()* flushes and closes standard I/O  
 48738 channels, thereby damaging the parent process' standard I/O data structures. (Even with *fork()*,  
 48739 it is wrong to call *exit()*, since buffered data would then be flushed twice.)

48740 If signal handlers are invoked in the child process after *vfork()*, they must follow the same rules  
 48741 as other code in the child process.



48742 **RATIONALE**

48743           None.

48744 **FUTURE DIRECTIONS**

48745           This function may be withdrawn in a future version.

48746 **SEE ALSO**48747           *exec*, *exit()*, *fork()*, *wait()*, the Base Definitions volume of IEEE Std 1003.1-200x, <**unistd.h**>48748 **CHANGE HISTORY**

48749           First released in Issue 4, Version 2.

48750 **Issue 5**

48751           Moved from X/OPEN UNIX extension to BASE.

48752 **Issue 6**

48753           Marked obsolescent.

48754 **NAME**

48755 vfprintf, vprintf, vsnprintf, vsprintf — format output of a stdarg argument list

48756 **SYNOPSIS**

48757 #include &lt;stdarg.h&gt;

48758 #include &lt;stdio.h&gt;

48759 int vfprintf(FILE \*restrict stream, const char \*restrict format,  
48760 va\_list ap);

48761 int vprintf(const char \*restrict format, va\_list ap);

48762 int vsnprintf(char \*restrict s, size\_t n, const char \*restrict format,  
48763 va\_list ap);

48764 int vsprintf(char \*restrict s, const char \*restrict format, va\_list ap);

48765 **DESCRIPTION**48766 **CX** The functionality described on this reference page is aligned with the ISO C standard. Any  
48767 conflict between the requirements described here and the ISO C standard is unintentional. This  
48768 volume of IEEE Std 1003.1-200x defers to the ISO C standard.48769 The *vprintf()*, *vfprintf()*, *vsnprintf()*, and *vsprintf()* functions shall be equivalent to *printf()*,  
48770 *fprintf()*, *snprintf()*, and *sprintf()* respectively, except that instead of being called with a variable  
48771 number of arguments, they are called with an argument list as defined by <stdarg.h>.48772 These functions shall not invoke the *va\_end* macro. As these functions invoke the *va\_arg* macro, |  
48773 the value of *ap* after the return is unspecified. |48774 **RETURN VALUE**48775 Refer to *fprintf()*.48776 **ERRORS**48777 Refer to *fprintf()*.48778 **EXAMPLES**

48779 None.

48780 **APPLICATION USAGE**48781 Applications using these functions should call *va\_end(ap)* afterwards to clean up.48782 **RATIONALE**

48783 None.

48784 **FUTURE DIRECTIONS**

48785 None.

48786 **SEE ALSO**48787 *fprintf()*, the Base Definitions volume of IEEE Std 1003.1-200x, <stdarg.h>, <stdio.h>48788 **CHANGE HISTORY**

48789 First released in Issue 1. Derived from Issue 1 of the SVID.

48790 **Issue 5**48791 The *vsnprintf()* function is added.48792 **Issue 6**48793 The *vfprintf()*, *vprintf()*, *vsnprintf()*, and *vsprintf()* functions are updated for alignment with the  
48794 ISO/IEC 9899:1999 standard.

48795 **NAME**48796 `vfscanf`, `vscanf`, `vsscanf` — format input of a stdarg list48797 **SYNOPSIS**48798 `#include <stdarg.h>`48799 `#include <stdio.h>`48800 `int vfscanf(FILE *restrict stream, const char *restrict format,`  
48801 `va_list arg);`48802 `int vscanf(const char *restrict format, va_list arg);`48803 `int vsscanf(const char *restrict s, const char *restrict format,`  
48804 `va_list arg);`48805 **DESCRIPTION**48806 **CX** The functionality described on this reference page is aligned with the ISO C standard. Any  
48807 conflict between the requirements described here and the ISO C standard is unintentional. This  
48808 volume of IEEE Std 1003.1-200x defers to the ISO C standard.48809 The `vscanf()`, `vfscanf()`, and `vsscanf()` functions shall be equivalent to the `scanf()`, `fscanf()`, and  
48810 `sscanf()` functions, respectively, except that instead of being called with a variable number of  
48811 arguments, they are called with an argument list as defined in the `<stdarg.h>` header. These  
48812 functions shall not invoke the `va_end` macro. As these functions invoke the `va_arg` macro, the  
48813 value of `ap` after the return is unspecified. |48814 **RETURN VALUE**48815 Refer to `fscanf()`.48816 **ERRORS**48817 Refer to `fscanf()`.48818 **EXAMPLES**

48819 None.

48820 **APPLICATION USAGE**48821 Applications using these functions should call `va_end(ap)` afterwards to clean up.48822 **RATIONALE**

48823 None.

48824 **FUTURE DIRECTIONS**

48825 None.

48826 **SEE ALSO**48827 `fscanf()`, the Base Definitions volume of IEEE Std 1003.1-200x, `<stdarg.h>`, `<stdio.h>`48828 **CHANGE HISTORY**

48829 First released in Issue 6. Derived from the ISO/IEC 9899:1999 standard.

48830 **NAME**

48831 vfwprintf, vswprintf, vwprintf — wide-character formatted output of a stdarg argument list

48832 **SYNOPSIS**

48833 #include &lt;stdarg.h&gt;

48834 #include &lt;stdio.h&gt;

48835 #include &lt;wchar.h&gt;

48836 int vfwprintf(FILE \*restrict stream, const wchar\_t \*restrict format,  
48837 va\_list arg);48838 int vswprintf(wchar\_t \*restrict ws, size\_t n,  
48839 const wchar\_t \*restrict format, va\_list arg);

48840 int vwprintf(const wchar\_t \*restrict format, va\_list arg);

48841 **DESCRIPTION**48842 **CX** The functionality described on this reference page is aligned with the ISO C standard. Any  
48843 conflict between the requirements described here and the ISO C standard is unintentional. This  
48844 volume of IEEE Std 1003.1-200x defers to the ISO C standard.48845 The *vfwprintf()*, *vswprintf()*, and *vwprintf()* functions shall be equivalent to *fwprintf()*, *swprintf()*,  
48846 and *wprintf()* respectively, except that instead of being called with a variable number of  
48847 arguments, they are called with an argument list as defined by <stdarg.h>.48848 These functions shall not invoke the *va\_end* macro. However, as these functions do invoke the  
48849 *va\_arg* macro, the value of *ap* after the return is unspecified. |48850 **RETURN VALUE**48851 Refer to *fwprintf()*.48852 **ERRORS**48853 Refer to *fwprintf()*.48854 **EXAMPLES**

48855 None.

48856 **APPLICATION USAGE**48857 Applications using these functions should call *va\_end(ap)* afterwards to clean up.48858 **RATIONALE**

48859 None.

48860 **FUTURE DIRECTIONS**

48861 None.

48862 **SEE ALSO**48863 *fwprintf()*, the Base Definitions volume of IEEE Std 1003.1-200x, <stdarg.h>, <stdio.h>,  
48864 <wchar.h>48865 **CHANGE HISTORY**48866 First released in Issue 5. Included for alignment with ISO/IEC 9899:1990/Amendment 1:1995  
48867 (E).48868 **Issue 6**48869 The *vfwprintf()*, *vswprintf()*, and *vwprintf()* prototypes are updated for alignment with the  
48870 ISO/IEC 9899:1999 standard. ()

48871 **NAME**

48872 vfwscanf, vswscanf, vwscanf — wide-character formatted input of a stdarg list

48873 **SYNOPSIS**

48874 #include &lt;stdarg.h&gt;

48875 #include &lt;stdio.h&gt;

48876 #include &lt;wchar.h&gt;

48877 int vfwscanf(FILE \*restrict stream, const wchar\_t \*restrict format,  
48878 va\_list arg);48879 int vswscanf(const wchar\_t \*restrict ws, const wchar\_t \*restrict format,  
48880 va\_list arg);

48881 int vwscanf(const wchar\_t \*restrict format, va\_list arg);

48882 **DESCRIPTION**48883 cx The functionality described on this reference page is aligned with the ISO C standard. Any  
48884 conflict between the requirements described here and the ISO C standard is unintentional. This  
48885 volume of IEEE Std 1003.1-200x defers to the ISO C standard.48886 The *vfwscanf()*, *vswscanf()*, and *vwscanf()* functions shall be equivalent to the *fwscanf()*, |  
48887 *swscanf()*, and *wscanf()* functions, respectively, except that instead of being called with a |  
48888 variable number of arguments, they are called with an argument list as defined in the <stdarg.h> |  
48889 header. These functions shall not invoke the *va\_end* macro. As these functions invoke the *va\_arg* |  
48890 macro, the value of *ap* after the return is unspecified. |48891 **RETURN VALUE**48892 Refer to *fwscanf()*.48893 **ERRORS**48894 Refer to *fwscanf()*.48895 **EXAMPLES**

48896 None.

48897 **APPLICATION USAGE**48898 Applications using these functions should call *va\_end(ap)* afterwards to clean up.48899 **RATIONALE**

48900 None.

48901 **FUTURE DIRECTIONS**

48902 None.

48903 **SEE ALSO**48904 *fwscanf()*, the Base Definitions volume of IEEE Std 1003.1-200x, <stdarg.h>, <stdio.h>,  
48905 <wchar.h>48906 **CHANGE HISTORY**

48907 First released in Issue 6. Derived from the ISO/IEC 9899:1999 standard.

48908 **NAME**

48909           vprintf — format output of a stdarg argument list

48910 **SYNOPSIS**

48911           #include <stdarg.h>

48912           #include <stdio.h>

48913           int vprintf(const char \*restrict *format*, va\_list *ap*);

48914 **DESCRIPTION**

48915           Refer to *fprintf()*.

48916 **NAME**

48917           vscanf — format input of a stdarg list

48918 **SYNOPSIS**

48919           #include &lt;stdarg.h&gt;

48920           #include &lt;stdio.h&gt;

48921           int vscanf(const char \*restrict *format*, va\_list *arg*);48922 **DESCRIPTION**48923           Refer to *vscanf()*.

48924 **NAME**

48925       vsprintf, vsprintf — format output of a stdarg argument list

48926 **SYNOPSIS**

48927       #include <stdarg.h>

48928       #include <stdio.h>

48929       int vsnprintf(char \*restrict *s*, size\_t *n*,

48930               const char \*restrict *format*, va\_list *ap*);

48931       int vsprintf(char \*restrict *s*, const char \*restrict *format*,

48932               va\_list *ap*);

48933 **DESCRIPTION**

48934       Refer to *fprintf()*.



48935 **NAME**

48936       vsscanf — format input of a stdarg list

48937 **SYNOPSIS**

48938       #include &lt;stdarg.h&gt;

48939       #include &lt;stdio.h&gt;

48940       int vsscanf(const char \*restrict *s*, const char \*restrict *format*,48941               va\_list *arg*);48942 **DESCRIPTION**48943       Refer to *vfscanf()*.

48944 **NAME**

48945       vswprintf — wide-character formatted output of a stdarg argument list

48946 **SYNOPSIS**

48947       #include <stdarg.h>

48948       #include <stdio.h>

48949       #include <wchar.h>

48950       int vswprintf(wchar\_t \*restrict ws, size\_t n,  
48951                    const wchar\_t \*restrict format, va\_list arg);

48952 **DESCRIPTION**

48953       Refer to *vfwprintf()*.

48954 **NAME**

48955 vswscanf — wide-character formatted input of a stdarg list

48956 **SYNOPSIS**

48957 #include &lt;stdarg.h&gt;

48958 #include &lt;stdio.h&gt;

48959 #include &lt;wchar.h&gt;

48960 int vswscanf(const wchar\_t \*restrict ws, const wchar\_t \*restrict format,  
48961 va\_list arg);48962 **DESCRIPTION**48963 Refer to *vfwscanf()*.

48964 **NAME**

48965       vwprintf — wide-character formatted output of a stdarg argument list

48966 **SYNOPSIS**

48967       #include <stdarg.h>

48968       #include <stdio.h>

48969       #include <wchar.h>

48970       int vwprintf(const wchar\_t \*restrict *format*, va\_list *arg*);

48971 **DESCRIPTION**

48972       Refer to *vfwprintf()*.

48973 **NAME**

48974 vwscanf — wide-character formatted input of a stdarg list

48975 **SYNOPSIS**

48976 #include &lt;stdarg.h&gt;

48977 #include &lt;stdio.h&gt;

48978 #include &lt;wchar.h&gt;

48979 int vwscanf(const wchar\_t \*restrict *format*, va\_list *arg*);48980 **DESCRIPTION**48981 Refer to *vfwscanf()*.

## 48982 NAME

48983 wait, waitpid — wait for a child process to stop or terminate

## 48984 SYNOPSIS

48985 #include &lt;sys/wait.h&gt;

48986 pid\_t wait(int \*stat\_loc);

48987 pid\_t waitpid(pid\_t pid, int \*stat\_loc, int options);

## 48988 DESCRIPTION

48989 The *wait()* and *waitpid()* functions shall obtain status information pertaining to one of the |  
 48990 caller's child processes. Various options permit status information to be obtained for child |  
 48991 processes that have terminated or stopped. If status information is available for two or more |  
 48992 child processes, the order in which their status is reported is unspecified. |

48993 The *wait()* function shall suspend execution of the calling thread until status information for one |  
 48994 of the terminated child processes of the calling process is available, or until delivery of a signal |  
 48995 whose action is either to execute a signal-catching function or to terminate the process. If more |  
 48996 than one thread is suspended in *wait()* or *waitpid()* awaiting termination of the same process, |  
 48997 exactly one thread shall return the process status at the time of the target process termination. If |  
 48998 status information is available prior to the call to *wait()*, return shall be immediate.

48999 The *waitpid()* function shall be equivalent to *wait()* if the *pid* argument is (**pid\_t**)−1 and the |  
 49000 *options* argument is 0. Otherwise, its behavior shall be modified by the values of the *pid* and |  
 49001 *options* arguments.

49002 The *pid* argument specifies a set of child processes for which *status* is requested. The *waitpid()* |  
 49003 function shall only return the status of a child process from this set:

- 49004 • If *pid* is equal to (**pid\_t**)−1, *status* is requested for any child process. In this respect, *waitpid()* |  
 49005 is then equivalent to *wait()*.
- 49006 • If *pid* is greater than 0, it specifies the process ID of a single child process for which *status* is |  
 49007 requested.
- 49008 • If *pid* is 0, *status* is requested for any child process whose process group ID is equal to that of |  
 49009 the calling process.
- 49010 • If *pid* is less than (**pid\_t**)−1, *status* is requested for any child process whose process group ID |  
 49011 is equal to the absolute value of *pid*.

49012 The *options* argument is constructed from the bitwise-inclusive OR of zero or more of the |  
 49013 following flags, defined in the <sys/wait.h> header:

49014 XSI WCONTINUED The *waitpid()* function shall report the status of any continued child process |  
 49015 specified by *pid* whose status has not been reported since it continued from a |  
 49016 job control stop.

49017 WNOHANG The *waitpid()* function shall not suspend execution of the calling thread if |  
 49018 *status* is not immediately available for one of the child processes specified by |  
 49019 *pid*.

49020 WUNTRACED The status of any child processes specified by *pid* that are stopped, and whose |  
 49021 status has not yet been reported since they stopped, shall also be reported to |  
 49022 the requesting process.

49023 XSI If the calling process has SA\_NOCLDWAIT set or has SIGCHLD set to SIG\_IGN, and the |  
 49024 process has no unwaited-for children that were transformed into zombie processes, the calling |  
 49025 thread shall block until all of the children of the process containing the calling thread terminate, |  
 49026 and *wait()* and *waitpid()* shall fail and set *errno* to [ECHILD].

49027 If *wait()* or *waitpid()* return because the status of a child process is available, these functions  
 49028 shall return a value equal to the process ID of the child process. In this case, if the value of the  
 49029 argument *stat\_loc* is not a null pointer, information shall be stored in the location pointed to by  
 49030 *stat\_loc*. The value stored at the location pointed to by *stat\_loc* shall be 0 if and only if the status  
 49031 returned is from a terminated child process that terminated by one of the following means:

- 49032 1. The process returned 0 from *main()*.
- 49033 2. The process called *\_exit()* or *exit()* with a *status* argument of 0.
- 49034 3. The process was terminated because the last thread in the process terminated.

49035 Regardless of its value, this information may be interpreted using the following macros, which  
 49036 are defined in `<sys/wait.h>` and evaluate to integral expressions; the *stat\_val* argument is the  
 49037 integer value pointed to by *stat\_loc*.

49038 **WIFEXITED(*stat\_val*)**

49039 Evaluates to a non-zero value if *status* was returned for a child process that terminated  
 49040 normally.

49041 **WEXITSTATUS(*stat\_val*)**

49042 If the value of **WIFEXITED(*stat\_val*)** is non-zero, this macro evaluates to the low-order 8 bits  
 49043 of the *status* argument that the child process passed to *\_exit()* or *exit()*, or the value the child  
 49044 process returned from *main()*.

49045 **WIFSIGNALED(*stat\_val*)**

49046 Evaluates to non-zero value if *status* was returned for a child process that terminated due to  
 49047 the receipt of a signal that was not caught (see `<signal.h>`).

49048 **WTERMSIG(*stat\_val*)**

49049 If the value of **WIFSIGNALED(*stat\_val*)** is non-zero, this macro evaluates to the number of  
 49050 the signal that caused the termination of the child process.

49051 **WIFSTOPPED(*stat\_val*)**

49052 Evaluates to a non-zero value if *status* was returned for a child process that is currently  
 49053 stopped.

49054 **WSTOPSIG(*stat\_val*)**

49055 If the value of **WIFSTOPPED(*stat\_val*)** is non-zero, this macro evaluates to the number of the  
 49056 signal that caused the child process to stop.

49057 XSI **WIFCONTINUED(*stat\_val*)**

49058 Evaluates to a non-zero value if *status* was returned for a child process that has continued  
 49059 from a job control stop.

49060 SPN It is unspecified whether the *status* value returned by calls to *wait()* or *waitpid()* for processes  
 49061 created by *posix\_spawn()* or *posix\_spawnnp()* can indicate a **WIFSTOPPED(*stat\_val*)** before  
 49062 subsequent calls to *wait()* or *waitpid()* indicate **WIFEXITED(*stat\_val*)** as the result of an error  
 49063 detected before the new process image starts executing.

49064 It is unspecified whether the *status* value returned by calls to *wait()* or *waitpid()* for processes  
 49065 created by *posix\_spawn()* or *posix\_spawnnp()* can indicate a **WIFSIGNALED(*stat\_val*)** if a signal is  
 49066 sent to the parent's process group after *posix\_spawn()* or *posix\_spawnnp()* is called.

49067 If the information pointed to by *stat\_loc* was stored by a call to *waitpid()* that specified the  
 49068 XSI **WUNTRACED** flag and did not specify the **WCONTINUED** flag, exactly one of the macros  
 49069 **WIFEXITED(\**stat\_loc*)**, **WIFSIGNALED(\**stat\_loc*)**, and **WIFSTOPPED(\**stat\_loc*)** shall evaluate to  
 49070 a non-zero value.

49071 If the information pointed to by *stat\_loc* was stored by a call to *waitpid()* that specified the  
 49072 XSI WUNTRACED and WCONTINUED flags, exactly one of the macros WIFEXITED(\**stat\_loc*),  
 49073 XSI WIFSIGNALED(\**stat\_loc*), WIFSTOPPED(\**stat\_loc*), and WIFCONTINUED(\**stat\_loc*) shall  
 49074 evaluate to a non-zero value.

49075 If the information pointed to by *stat\_loc* was stored by a call to *waitpid()* that did not specify the  
 49076 XSI WUNTRACED or WCONTINUED flags, or by a call to the *wait()* function, exactly one of the  
 49077 macros WIFEXITED(\**stat\_loc*) and WIFSIGNALED(\**stat\_loc*) shall evaluate to a non-zero value.

49078 If the information pointed to by *stat\_loc* was stored by a call to *waitpid()* that did not specify the  
 49079 XSI WUNTRACED flag and specified the WCONTINUED flag, or by a call to the *wait()* function,  
 49080 XSI exactly one of the macros WIFEXITED(\**stat\_loc*), WIFSIGNALED(\**stat\_loc*), and  
 49081 WIFCONTINUED(\**stat\_loc*) shall evaluate to a non-zero value.

49082 If `_POSIX_REALTIME_SIGNALS` is defined, and the implementation queues the SIGCHLD  
 49083 signal, then if *wait()* or *waitpid()* returns because the status of a child process is available, any  
 49084 pending SIGCHLD signal associated with the process ID of the child process shall be discarded.  
 49085 Any other pending SIGCHLD signals shall remain pending.

49086 Otherwise, if SIGCHLD is blocked, if *wait()* or *waitpid()* return because the status of a child  
 49087 process is available, any pending SIGCHLD signal shall be cleared unless the status of another  
 49088 child process is available.

49089 For all other conditions, it is unspecified whether child *status* will be available when a SIGCHLD  
 49090 signal is delivered.

49091 There may be additional implementation-defined circumstances under which *wait()* or *waitpid()*  
 49092 report *status*. This shall not occur unless the calling process or one of its child processes explicitly  
 49093 makes use of a non-standard extension. In these cases the interpretation of the reported *status* is  
 49094 implementation-defined.

49095 XSI If a parent process terminates without waiting for all of its child processes to terminate, the  
 49096 remaining child processes shall be assigned a new parent process ID corresponding to an  
 49097 implementation-defined system process.

#### 49098 RETURN VALUE

49099 If *wait()* or *waitpid()* returns because the status of a child process is available, these functions  
 49100 shall return a value equal to the process ID of the child process for which *status* is reported. If  
 49101 *wait()* or *waitpid()* returns due to the delivery of a signal to the calling process, `-1` shall be  
 49102 returned and *errno* set to `[EINTR]`. If *waitpid()* was invoked with `WNOHANG` set in *options*, it  
 49103 has at least one child process specified by *pid* for which *status* is not available, and *status* is not  
 49104 available for any process specified by *pid*, `0` is returned. Otherwise, `(pid_t)-1` shall be returned,  
 49105 and *errno* set to indicate the error.

#### 49106 ERRORS

49107 The *wait()* function shall fail if:

49108 `[ECHILD]` The calling process has no existing unwaited-for child processes.

49109 `[EINTR]` The function was interrupted by a signal. The value of the location pointed to  
 49110 by *stat\_loc* is undefined.

49111 The *waitpid()* function shall fail if:

49112 `[ECHILD]` The process specified by *pid* does not exist or is not a child of the calling  
 49113 process, or the process group specified by *pid* does not exist or does not have  
 49114 any member process that is a child of the calling process.



49115 [EINTR] The function was interrupted by a signal. The value of the location pointed to  
49116 by *stat\_loc* is undefined.

49117 [EINVAL] The *options* argument is not valid.

#### 49118 EXAMPLES

49119 None.

#### 49120 APPLICATION USAGE

49121 None.

#### 49122 RATIONALE

49123 A call to the *wait()* or *waitpid()* function only returns *status* on an immediate child process of the  
49124 calling process; that is, a child that was produced by a single *fork()* call (perhaps followed by an  
49125 *exec* or other function calls) from the parent. If a child produces grandchildren by further use of  
49126 *fork()*, none of those grandchildren nor any of their descendants affect the behavior of a *wait()*  
49127 from the original parent process. Nothing in this volume of IEEE Std 1003.1-200x prevents an  
49128 implementation from providing extensions that permit a process to get *status* from a grandchild  
49129 or any other process, but a process that does not use such extensions must be guaranteed to see  
49130 *status* from only its direct children.

49131 The *waitpid()* function is provided for three reasons:

- 49132 1. To support job control
- 49133 2. To permit a non-blocking version of the *wait()* function
- 49134 3. To permit a library routine, such as *system()* or *pclose()*, to wait for its children without  
49135 interfering with other terminated children for which the process has not waited

49136 The first two of these facilities are based on the *wait3()* function provided by 4.3 BSD. The  
49137 function uses the *options* argument, which is equivalent to an argument to *wait3()*. The  
49138 WUNTRACED flag is used only in conjunction with job control on systems supporting job  
49139 control. Its name comes from 4.3 BSD and refers to the fact that there are two types of stopped  
49140 processes in that implementation: processes being traced via the *ptrace()* debugging facility and  
49141 (untraced) processes stopped by job control signals. Since *ptrace()* is not part of this volume of  
49142 IEEE Std 1003.1-200x, only the second type is relevant. The name WUNTRACED was retained  
49143 because its usage is the same, even though the name is not intuitively meaningful in this context.

49144 The third reason for the *waitpid()* function is to permit independent sections of a process to  
49145 spawn and wait for children without interfering with each other. For example, the following  
49146 problem occurs in developing a portable shell, or command interpreter:

```
49147 stream = popen("/bin/true");
49148 (void) system("sleep 100");
49149 (void) pclose(stream);
```

49150 On all historical implementations, the final *pclose()* fails to reap the *wait()* *status* of the *popen()*.

49151 The status values are retrieved by macros, rather than given as specific bit encodings as they are  
49152 in most historical implementations (and thus expected by existing programs). This was  
49153 necessary to eliminate a limitation on the number of signals an implementation can support that  
49154 was inherent in the traditional encodings. This volume of IEEE Std 1003.1-200x does require that  
49155 a *status* value of zero corresponds to a process calling *\_exit(0)*, as this is the most common  
49156 encoding expected by existing programs. Some of the macro names were adopted from 4.3 BSD.

49157 These macros syntactically operate on an arbitrary integer value. The behavior is undefined  
49158 unless that value is one stored by a successful call to *wait()* or *waitpid()* in the location pointed  
49159 to by the *stat\_loc* argument. An early proposal attempted to make this clearer by specifying each

49160 argument as *\*stat\_loc* rather than *stat\_val*. However, that did not follow the conventions of other  
49161 specifications in this volume of IEEE Std 1003.1-200x or traditional usage. It also could have  
49162 implied that the argument to the macro must literally be *\*stat\_loc*; in fact, that value can be  
49163 stored or passed as an argument to other functions before being interpreted by these macros.

49164 The extension that affects *wait()* and *waitpid()* and is common in historical implementations is  
49165 the *ptrace()* function. It is called by a child process and causes that child to stop and return a  
49166 *status* that appears identical to the *status* indicated by WIFSTOPPED. The *status* of *ptrace()*  
49167 children is traditionally returned regardless of the WUNTRACED flag (or by the *wait()*  
49168 function). Most applications do not need to concern themselves with such extensions because  
49169 they have control over what extensions they or their children use. However, applications, such  
49170 as command interpreters, that invoke arbitrary processes may see this behavior when those  
49171 arbitrary processes misuse such extensions.

49172 Implementations that support *core* file creation or other implementation-defined actions on  
49173 termination of some processes traditionally provide a bit in the *status* returned by *wait()* to  
49174 indicate that such actions have occurred.

49175 Allowing the *wait()* family of functions to discard a pending SIGCHLD signal that is associated  
49176 with a successfully waited-for child process puts them into the *sigwait()* and *sigwaitinfo()*  
49177 category with respect to SIGCHLD.

49178 This definition allows implementations to treat a pending SIGCHLD signal as accepted by the  
49179 process in *wait()*, with the same meaning of “accepted” as when that word is applied to the  
49180 *sigwait()* family of functions.

49181 Allowing the *wait()* family of functions to behave this way permits an implementation to be able  
49182 to deal precisely with SIGCHLD signals.

49183 In particular, an implementation that does accept (discard) the SIGCHLD signal can make the  
49184 following guarantees regardless of the queuing depth of signals in general (the list of waitable  
49185 children can hold the SIGCHLD queue):

- 49186 1. If a SIGCHLD signal handler is established via *sigaction()* without the SA\_RESETHAND  
49187 flag, SIGCHLD signals can be accurately counted; that is, exactly one SIGCHLD signal will  
49188 be delivered to or accepted by the process for every child process that terminates.
- 49189 2. A single *wait()* issued from a SIGCHLD signal handler can be guaranteed to return  
49190 immediately with status information for a child process.
- 49191 3. When SA\_SIGINFO is requested, the SIGCHLD signal handler can be guaranteed to  
49192 receive a non-NULL pointer to a **siginfo\_t** structure that describes a child process for  
49193 which a wait via *waitpid()* or *waitid()* will not block or fail.
- 49194 4. The *system()* function will not cause a process's SIGCHLD handler to be called as a result of  
49195 the *fork()/exec* executed within *system()* because *system()* will accept the SIGCHLD signal  
49196 when it performs a *waitpid()* for its child process. This is a desirable behavior of *system()*  
49197 so that it can be used in a library without causing side effects to the application linked with  
49198 the library.

49199 An implementation that does not permit the *wait()* family of functions to accept (discard) a  
49200 pending SIGCHLD signal associated with a successfully waited-for child, cannot make the  
49201 guarantees described above for the following reasons:

#### 49202 Guarantee #1

49203 Although it might be assumed that reliable queuing of all SIGCHLD signals generated by  
49204 the system can make this guarantee, the counter example is the case of a process that blocks  
49205 SIGCHLD and performs an indefinite loop of *fork()/wait()* operations. If the

49206 implementation supports queued signals, then eventually the system will run out of  
 49207 memory for the queue. The guarantee cannot be made because there must be some limit to  
 49208 the depth of queuing.

49209 Guarantees #2 and #3

49210 These cannot be guaranteed unless the *wait()* family of functions accepts the SIGCHLD  
 49211 signal. Otherwise, a *fork()/wait()* executed while SIGCHLD is blocked (as in the *system()*  
 49212 function) will result in an invocation of the handler when SIGCHLD is unblocked, after the  
 49213 process has disappeared.

49214 Guarantee #4

49215 Although possible to make this guarantee, *system()* would have to set the SIGCHLD  
 49216 handler to SIG\_DFL so that the SIGCHLD signal generated by its *fork()* would be discarded  
 49217 (the SIGCHLD default action is to be ignored), then restore it to its previous setting. This  
 49218 would have the undesirable side effect of discarding all SIGCHLD signals pending to the  
 49219 process.

#### 49220 FUTURE DIRECTIONS

49221 None.

#### 49222 SEE ALSO

49223 *exec*, *exit()*, *fork()*, *waitid()*, the Base Definitions volume of IEEE Std 1003.1-200x, <sys/types.h>,  
 49224 <sys/wait.h>

#### 49225 CHANGE HISTORY

49226 First released in Issue 1. Derived from Issue 1 of the SVID.

#### 49227 Issue 5

49228 The DESCRIPTION is updated for alignment with the POSIX Threads Extension.

#### 49229 Issue 6

49230 In the SYNOPSIS, the optional include of the <sys/types.h> header is removed.

49231 The following new requirements on POSIX implementations derive from alignment with the  
 49232 Single UNIX Specification:

- 49233 • The requirement to include <sys/types.h> has been removed. Although <sys/types.h> was  
 49234 required for conforming implementations of previous POSIX specifications, it was not  
 49235 required for UNIX applications.

49236 The following changes were made to align with the IEEE P1003.1a draft standard:

- 49237 • The processing of the SIGCHLD signal and the [ECHILD] error is clarified.

49238 The semantics of WIFSTOPPED(*stat\_val*), WIFEXITED(*stat\_val*), and WIFSIGNALED(*stat\_val*)  
 49239 are defined with respect to *posix\_spawn()* or *posix\_spawnnp()* for alignment with  
 49240 IEEE Std 1003.1d-1999.

49241 The DESCRIPTION is updated for alignment with the ISO/IEC 9899:1999 standard.

## 49242 NAME

49243 waitid — wait for a child process to change state

## 49244 SYNOPSIS

49245 XSI #include &lt;sys/wait.h&gt;

49246 int waitid(idtype\_t idtype, id\_t id, siginfo\_t \*infop, int options);

49247

## 49248 DESCRIPTION

49249 The *waitid()* function shall suspend the calling thread until one child of the process containing  
 49250 the calling thread changes state. It records the current state of a child in the structure pointed to  
 49251 by *infop*. If a child process changed state prior to the call to *waitid()*, *waitid()* shall return  
 49252 immediately. If more than one thread is suspended in *wait()* or *waitpid()* waiting termination of  
 49253 the same process, exactly one thread shall return the process status at the time of the target  
 49254 process termination.

49255 The *idtype* and *id* arguments are used to specify which children *waitid()* waits for.

49256 If *idtype* is P\_PID, *waitid()* shall wait for the child with a process ID equal to (**pid\_t**)*id*.

49257 If *idtype* is P\_PGID, *waitid()* shall wait for any child with a process group ID equal to (**pid\_t**)*id*.

49258 If *idtype* is P\_ALL, *waitid()* shall wait for any children and *id* is ignored.

49259 The *options* argument is used to specify which state changes *waitid()* shall wait for. It is formed  
 49260 by OR'ing together one or more of the following flags:

49261 WEXITED Wait for processes that have exited.

49262 WSTOPPED Status shall be returned for any child that has stopped upon receipt of a signal.

49263 WCONTINUED Status shall be returned for any child that was stopped and has been  
 49264 continued.

49265 WNOHANG Return immediately if there are no children to wait for.

49266 WNOWAIT Keep the process whose status is returned in *infop* in a waitable state. This  
 49267 shall not affect the state of the process; the process may be waited for again  
 49268 after this call completes.

49269 The application shall ensure that the *infop* argument points to a **siginfo\_t** structure. If *waitid()*  
 49270 returns because a child process was found that satisfied the conditions indicated by the  
 49271 arguments *idtype* and *options*, then the structure pointed to by *infop* shall be filled in by the  
 49272 system with the status of the process. The *si\_signo* member shall always be equal to SIGCHLD.

## 49273 RETURN VALUE

49274 If WNOHANG was specified and there are no children to wait for, 0 shall be returned. If *waitid()* |  
 49275 returns due to the change of state of one of its children, 0 shall be returned. Otherwise, -1 shall  
 49276 be returned and *errno* set to indicate the error.

## 49277 ERRORS

49278 The *waitid()* function shall fail if:

49279 [ECHILD] The calling process has no existing unwaited-for child processes.

49280 [EINTR] The *waitid()* function was interrupted by a signal.

49281 [EINVAL] An invalid value was specified for *options*, or *idtype* and *id* specify an invalid  
 49282 set of processes.

49283 **EXAMPLES**

49284 None.

49285 **APPLICATION USAGE**

49286 None.

49287 **RATIONALE**

49288 None.

49289 **FUTURE DIRECTIONS**

49290 None.

49291 **SEE ALSO**49292 *exec*, *exit()*, *wait()*, the Base Definitions volume of IEEE Std 1003.1-200x, <sys/wait.h>49293 **CHANGE HISTORY**

49294 First released in Issue 4, Version 2.

49295 **Issue 5**

49296 Moved from X/OPEN UNIX extension to BASE.

49297 The DESCRIPTION is updated for alignment with the POSIX Threads Extension.

49298 **Issue 6**

49299 The DESCRIPTION is updated to avoid use of the term “must” for application requirements.

49300 **NAME**

49301           waitpid — wait for a child process to stop or terminate

49302 **SYNOPSIS**

49303           #include <sys/wait.h>

49304           pid\_t waitpid(pid\_t *pid*, int *\*stat\_loc*, int *options*);

49305 **DESCRIPTION**

49306           Refer to *wait()*.

49307 **NAME**49308 `wrtomb` — convert a wide-character code to a character (restartable)49309 **SYNOPSIS**49310 `#include <stdio.h>`49311 `size_t wrtomb(char *restrict s, wchar_t wc, mbstate_t *restrict ps);`49312 **DESCRIPTION**

49313 cx The functionality described on this reference page is aligned with the ISO C standard. Any  
 49314 conflict between the requirements described here and the ISO C standard is unintentional. This  
 49315 volume of IEEE Std 1003.1-200x defers to the ISO C standard.

49316 If *s* is a null pointer, the `wrtomb()` function shall be equivalent to the call:49317 `wrtomb(buf, L'\0', ps)`49318 where *buf* is an internal buffer.

49319 If *s* is not a null pointer, the `wrtomb()` function shall determine the number of bytes needed to  
 49320 represent the character that corresponds to the wide character given by *wc* (including any shift  
 49321 sequences), and store the resulting bytes in the array whose first element is pointed to by *s*. At  
 49322 most {MB\_CUR\_MAX} bytes are stored. If *wc* is a null wide character, a null byte shall be stored,  
 49323 preceded by any shift sequence needed to restore the initial shift state. The resulting state  
 49324 described shall be the initial conversion state.

49325 If *ps* is a null pointer, the `wrtomb()` function shall use its own internal `mbstate_t` object, which is  
 49326 initialized at program start-up to the initial conversion state. Otherwise, the `mbstate_t` object  
 49327 pointed to by *ps* shall be used to completely describe the current conversion state of the  
 49328 associated character sequence. The implementation shall behave as if no function defined in this  
 49329 volume of IEEE Std 1003.1-200x calls `wrtomb()`.

49330 cx If the application uses any of the `_POSIX_THREAD_SAFE_FUNCTIONS` or `_POSIX_THREADS`  
 49331 functions, the application shall ensure that the `wrtomb()` function is called with a non-NULL *ps*  
 49332 argument.

49333 The behavior of this function shall be affected by the `LC_CTYPE` category of the current locale.49334 **RETURN VALUE**

49335 The `wrtomb()` function shall return the number of bytes stored in the array object (including any  
 49336 shift sequences). When *wc* is not a valid wide character, an encoding error shall occur. In this  
 49337 case, the function shall store the value of the macros [EILSEQ] in *errno* and shall return  
 49338 (`size_t`)-1; the conversion state shall be undefined.

49339 **ERRORS**49340 The `wrtomb()` function may fail if:49341 cx [EINVAL] *ps* points to an object that contains an invalid conversion state.

49342 [EILSEQ] Invalid wide-character code is detected.

49343 **EXAMPLES**

49344 None.

49345 **APPLICATION USAGE**

49346 None.

49347 **RATIONALE**

49348 None.

49349 **FUTURE DIRECTIONS**

49350 None.

49351 **SEE ALSO**49352 *mbstinit()*, the Base Definitions volume of IEEE Std 1003.1-200x, <**wchar.h**>49353 **CHANGE HISTORY**49354 First released in Issue 5. Included for alignment with ISO/IEC 9899:1990/Amendment 1:1995  
49355 (E).49356 **Issue 6**

49357 In the DESCRIPTION, a note on using this function in a threaded application is added.

49358 Extensions beyond the ISO C standard are now marked.

49359 The DESCRIPTION is updated to avoid use of the term “must” for application requirements.

49360 The *wcrtomb()* prototype is updated for alignment with the ISO/IEC 9899:1999 standard.



49361 **NAME**49362 `wscat` — concatenate two wide-character strings49363 **SYNOPSIS**49364 `#include <wchar.h>`49365 `wchar_t *wscat(wchar_t *restrict ws1, const wchar_t *restrict ws2);`49366 **DESCRIPTION**49367 **CX** The functionality described on this reference page is aligned with the ISO C standard. Any  
49368 conflict between the requirements described here and the ISO C standard is unintentional. This  
49369 volume of IEEE Std 1003.1-200x defers to the ISO C standard.49370 The `wscat()` function shall append a copy of the wide-character string pointed to by `ws2`  
49371 (including the terminating null wide-character code) to the end of the wide-character string  
49372 pointed to by `ws1`. The initial wide-character code of `ws2` shall overwrite the null wide-character  
49373 code at the end of `ws1`. If copying takes place between objects that overlap, the behavior is  
49374 undefined.49375 **RETURN VALUE**49376 The `wscat()` function shall return `ws1`; no return value is reserved to indicate an error.49377 **ERRORS**

49378 No errors are defined.

49379 **EXAMPLES**

49380 None.

49381 **APPLICATION USAGE**

49382 None.

49383 **RATIONALE**

49384 None.

49385 **FUTURE DIRECTIONS**

49386 None.

49387 **SEE ALSO**49388 `wscncat()`, the Base Definitions volume of IEEE Std 1003.1-200x, `<wchar.h>`49389 **CHANGE HISTORY**

49390 First released in Issue 4. Derived from the MSE working draft.

49391 **Issue 6**49392 The Open Group Corrigendum U040/2 is applied. In the RETURN VALUE section, `s1` is changed  
49393 to `ws1`.49394 The `wscat()` prototype is updated for alignment with the ISO/IEC 9899:1999 standard.

49395 **NAME**49396 `wchr` — wide-character string scanning operation49397 **SYNOPSIS**49398 `#include <wchar.h>`49399 `wchar_t *wchr(const wchar_t *ws, wchar_t wc);`49400 **DESCRIPTION**

49401 `cx` The functionality described on this reference page is aligned with the ISO C standard. Any  
49402 conflict between the requirements described here and the ISO C standard is unintentional. This  
49403 volume of IEEE Std 1003.1-200x defers to the ISO C standard.

49404 The `wchr()` function shall locate the first occurrence of `wc` in the wide-character string pointed  
49405 to by `ws`. The application shall ensure that the value of `wc` is a character representable as a type  
49406 `wchar_t` and a wide-character code corresponding to a valid character in the current locale. The  
49407 terminating null wide-character code is considered to be part of the wide-character string.

49408 **RETURN VALUE**

49409 Upon completion, `wchr()` shall return a pointer to the wide-character code, or a null pointer if  
49410 the wide-character code is not found.

49411 **ERRORS**

49412 No errors are defined.

49413 **EXAMPLES**

49414 None.

49415 **APPLICATION USAGE**

49416 None.

49417 **RATIONALE**

49418 None.

49419 **FUTURE DIRECTIONS**

49420 None.

49421 **SEE ALSO**49422 `wchr()`, the Base Definitions volume of IEEE Std 1003.1-200x, `<wchar.h>`49423 **CHANGE HISTORY**

49424 First released in Issue 4. Derived from the MSE working draft.

49425 **Issue 6**

49426 The DESCRIPTION is updated to avoid use of the term “must” for application requirements.

49427 **NAME**49428        `wcsncmp` — compare two wide-character strings49429 **SYNOPSIS**49430        `#include <wchar.h>`49431        `int wcsncmp(const wchar_t *ws1, const wchar_t *ws2);`49432 **DESCRIPTION**49433 **CX**        The functionality described on this reference page is aligned with the ISO C standard. Any  
49434        conflict between the requirements described here and the ISO C standard is unintentional. This  
49435        volume of IEEE Std 1003.1-200x defers to the ISO C standard.49436        The `wcsncmp()` function shall compare the wide-character string pointed to by `ws1` to the wide-  
49437        character string pointed to by `ws2`.49438        The sign of a non-zero return value shall be determined by the sign of the difference between the  
49439        values of the first pair of wide-character codes that differ in the objects being compared. |49440 **RETURN VALUE**49441        Upon completion, `wcsncmp()` shall return an integer greater than, equal to, or less than 0, if the  
49442        wide-character string pointed to by `ws1` is greater than, equal to, or less than the wide-character  
49443        string pointed to by `ws2`, respectively.49444 **ERRORS**

49445        No errors are defined.

49446 **EXAMPLES**

49447        None.

49448 **APPLICATION USAGE**

49449        None.

49450 **RATIONALE**

49451        None.

49452 **FUTURE DIRECTIONS**

49453        None.

49454 **SEE ALSO**49455        `wcsncmp()`, the Base Definitions volume of IEEE Std 1003.1-200x, `<wchar.h>`49456 **CHANGE HISTORY**

49457        First released in Issue 4. Derived from the MSE working draft.

49458 **NAME**

49459 wscoll — wide-character string comparison using collating information

49460 **SYNOPSIS**

49461 #include &lt;wchar.h&gt;

49462 int wscoll(const wchar\_t \*ws1, const wchar\_t \*ws2);

49463 **DESCRIPTION**49464 CX The functionality described on this reference page is aligned with the ISO C standard. Any  
49465 conflict between the requirements described here and the ISO C standard is unintentional. This  
49466 volume of IEEE Std 1003.1-200x defers to the ISO C standard.49467 The *wscoll()* function shall compare the wide-character string pointed to by *ws1* to the wide-  
49468 character string pointed to by *ws2*, both interpreted as appropriate to the *LC\_COLLATE* category  
49469 of the current locale.49470 CX The *wscoll()* function shall not change the setting of *errno* if successful.49471 An application wishing to check for error situations should set *errno* to 0 before calling *wscoll()*.  
49472 If *errno* is non-zero on return, an error has occurred.49473 **RETURN VALUE**49474 Upon successful completion, *wscoll()* shall return an integer greater than, equal to, or less than  
49475 0, according to whether the wide-character string pointed to by *ws1* is greater than, equal to, or  
49476 less than the wide-character string pointed to by *ws2*, when both are interpreted as appropriate  
49477 CX to the current locale. On error, *wscoll()* shall set *errno*, but no return value is reserved to  
49478 indicate an error.49479 **ERRORS**49480 The *wscoll()* function may fail if:49481 CX [EINVAL] The *ws1* or *ws2* arguments contain wide-character codes outside the domain of  
49482 the collating sequence.49483 **EXAMPLES**

49484 None.

49485 **APPLICATION USAGE**49486 The *wcsxfrm()* and *wscmp()* functions should be used for sorting large lists.49487 **RATIONALE**

49488 None.

49489 **FUTURE DIRECTIONS**

49490 None.

49491 **SEE ALSO**49492 *wscmp()*, *wcsxfrm()*, the Base Definitions volume of IEEE Std 1003.1-200x, <wchar.h>49493 **CHANGE HISTORY**

49494 First released in Issue 4. Derived from the MSE working draft.

49495 **Issue 5**

49496 Moved from ENHANCED I18N to BASE and the [ENOSYS] error is removed.

49497 The DESCRIPTION is updated to indicate that *errno* is not changed if the function is successful.

49498 **NAME**

49499           wcscpy — copy a wide-character string

49500 **SYNOPSIS**

49501           #include &lt;wchar.h&gt;

49502           wchar\_t \*wcscpy(wchar\_t \*restrict ws1, const wchar\_t \*restrict ws2);

49503 **DESCRIPTION**49504 cx       The functionality described on this reference page is aligned with the ISO C standard. Any  
49505       conflict between the requirements described here and the ISO C standard is unintentional. This  
49506       volume of IEEE Std 1003.1-200x defers to the ISO C standard.49507       The *wcscpy()* function shall copy the wide-character string pointed to by *ws2* (including the  
49508       terminating null wide-character code) into the array pointed to by *ws1*. If copying takes place  
49509       between objects that overlap, the behavior is undefined.49510 **RETURN VALUE**49511       The *wcscpy()* function shall return *ws1*; no return value is reserved to indicate an error.49512 **ERRORS**

49513       No errors are defined.

49514 **EXAMPLES**

49515       None.

49516 **APPLICATION USAGE**

49517       None.

49518 **RATIONALE**

49519       None.

49520 **FUTURE DIRECTIONS**

49521       None.

49522 **SEE ALSO**49523       *wscncpy()*, the Base Definitions volume of IEEE Std 1003.1-200x, <**wchar.h**>49524 **CHANGE HISTORY**

49525       First released in Issue 4. Derived from the MSE working draft.

49526 **Issue 6**49527       The *wcscpy()* prototype is updated for alignment with the ISO/IEC 9899:1999 standard.

49528 **NAME**

49529 wscspn — get length of a complementary wide substring

49530 **SYNOPSIS**

49531 #include <wchar.h>

49532 size\_t wscspn(const wchar\_t \*ws1, const wchar\_t \*ws2);

49533 **DESCRIPTION**

49534 cx The functionality described on this reference page is aligned with the ISO C standard. Any  
49535 conflict between the requirements described here and the ISO C standard is unintentional. This  
49536 volume of IEEE Std 1003.1-200x defers to the ISO C standard.

49537 The *wscspn()* function shall compute the length (in wide characters) of the maximum initial |  
49538 segment of the wide-character string pointed to by *ws1* which consists entirely of wide-character |  
49539 codes *not* from the wide-character string pointed to by *ws2*.

49540 **RETURN VALUE**

49541 The *wscspn()* function shall return the length of the initial substring of *ws1*; no return value is  
49542 reserved to indicate an error.

49543 **ERRORS**

49544 No errors are defined.

49545 **EXAMPLES**

49546 None.

49547 **APPLICATION USAGE**

49548 None.

49549 **RATIONALE**

49550 None.

49551 **FUTURE DIRECTIONS**

49552 None.

49553 **SEE ALSO**

49554 *wcspn()*, the Base Definitions volume of IEEE Std 1003.1-200x, <wchar.h>

49555 **CHANGE HISTORY**

49556 First released in Issue 4. Derived from the MSE working draft.

49557 **Issue 5**

49558 The RETURN VALUE section is updated to indicate that *wscspn()* returns the length of *ws1*,  
49559 rather than *ws1* itself.

49560 **NAME**49561 `wcsftime` — convert date and time to a wide-character string49562 **SYNOPSIS**49563 `#include <wchar.h>`49564 `size_t wcsftime(wchar_t *restrict wcs, size_t maxsize,`  
49565 `const wchar_t *restrict format, const struct tm *restrict timeptr);` |49566 **DESCRIPTION**49567 **CX** The functionality described on this reference page is aligned with the ISO C standard. Any  
49568 conflict between the requirements described here and the ISO C standard is unintentional. This  
49569 volume of IEEE Std 1003.1-200x defers to the ISO C standard.49570 The `wcsftime()` function shall be equivalent to the `strftime()` function, except that:

- 49571 • The argument `wcs` points to the initial element of an array of wide characters into which the  
49572 generated output is to be placed.
- 49573 • The argument `maxsize` indicates the maximum number of wide characters to be placed in the  
49574 output array.
- 49575 • The argument `format` is a wide-character string and the conversion specifications are replaced  
49576 by corresponding sequences of wide characters.
- 49577 • The return value indicates the number of wide characters placed in the output array.

49578 If copying takes place between objects that overlap, the behavior is undefined.

49579 **RETURN VALUE**49580 If the total number of resulting wide-character codes including the terminating null wide-  
49581 character code is no more than `maxsize`, `wcsftime()` shall return the number of wide-character  
49582 codes placed into the array pointed to by `wcs`, not including the terminating null wide-character  
49583 code. Otherwise, zero is returned and the contents of the array are unspecified. |49584 **ERRORS**

49585 No errors are defined.

49586 **EXAMPLES**

49587 None.

49588 **APPLICATION USAGE**

49589 None.

49590 **RATIONALE**

49591 None.

49592 **FUTURE DIRECTIONS**

49593 None.

49594 **SEE ALSO**49595 `strftime()`, the Base Definitions volume of IEEE Std 1003.1-200x, `<wchar.h>`49596 **CHANGE HISTORY**

49597 First released in Issue 4.

49598 **Issue 5**

49599 Moved from ENHANCED I18N to BASE and the [ENOSYS] error is removed.

49600 Aligned with ISO/IEC 9899:1990/Amendment 1:1995 (E). Specifically, the type of the `format`  
49601 argument is changed from `const char *` to `const wchar_t *`.

49602 **Issue 6**

49603

The *wcsftime()* prototype is updated for alignment with the ISO/IEC 9899:1999 standard.



49604 **NAME**

49605           wcslen — get wide-character string length

49606 **SYNOPSIS**

49607           #include &lt;wchar.h&gt;

49608           size\_t wcslen(const wchar\_t \*ws);

49609 **DESCRIPTION**

49610 cx       The functionality described on this reference page is aligned with the ISO C standard. Any  
49611       conflict between the requirements described here and the ISO C standard is unintentional. This  
49612       volume of IEEE Std 1003.1-200x defers to the ISO C standard.

49613       The *wcslen()* function shall compute the number of wide-character codes in the wide-character  
49614       string to which *ws* points, not including the terminating null wide-character code.

49615 **RETURN VALUE**

49616       The *wcslen()* function shall return the length of *ws*; no return value is reserved to indicate an  
49617       error.

49618 **ERRORS**

49619       No errors are defined.

49620 **EXAMPLES**

49621       None.

49622 **APPLICATION USAGE**

49623       None.

49624 **RATIONALE**

49625       None.

49626 **FUTURE DIRECTIONS**

49627       None.

49628 **SEE ALSO**49629       The Base Definitions volume of IEEE Std 1003.1-200x, <**wchar.h**>49630 **CHANGE HISTORY**

49631       First released in Issue 4. Derived from the MSE working draft.

49632 **NAME**

49633       wcsncat — concatenate a wide-character string with part of another

49634 **SYNOPSIS**

49635       #include &lt;wchar.h&gt;

49636       wchar\_t \*wcsncat(wchar\_t \*restrict ws1, const wchar\_t \*restrict ws2,  
49637                       size\_t n);49638 **DESCRIPTION**49639 **CX**       The functionality described on this reference page is aligned with the ISO C standard. Any  
49640       conflict between the requirements described here and the ISO C standard is unintentional. This  
49641       volume of IEEE Std 1003.1-200x defers to the ISO C standard.49642       The *wcsncat()* function shall append not more than *n* wide-character codes (a null wide-  
49643       character code and wide-character codes that follow it are not appended) from the array pointed  
49644       to by *ws2* to the end of the wide-character string pointed to by *ws1*. The initial wide-character  
49645       code of *ws2* shall overwrite the null wide-character code at the end of *ws1*. A terminating null  
49646       wide-character code shall always be appended to the result. If copying takes place between |  
49647       objects that overlap, the behavior is undefined.49648 **RETURN VALUE**49649       The *wcsncat()* function shall return *ws1*; no return value is reserved to indicate an error.49650 **ERRORS**

49651       No errors are defined.

49652 **EXAMPLES**

49653       None.

49654 **APPLICATION USAGE**

49655       None.

49656 **RATIONALE**

49657       None.

49658 **FUTURE DIRECTIONS**

49659       None.

49660 **SEE ALSO**49661       *wscat()*, the Base Definitions volume of IEEE Std 1003.1-200x, <**wchar.h**>49662 **CHANGE HISTORY**

49663       First released in Issue 4. Derived from the MSE working draft.

49664 **Issue 6**49665       The *wcsncat()* prototype is updated for alignment with the ISO/IEC 9899: 1999 standard.

49666 **NAME**49667 `wcsncmp` — compare part of two wide-character strings49668 **SYNOPSIS**49669 `#include <wchar.h>`49670 `int wcsncmp(const wchar_t *ws1, const wchar_t *ws2, size_t n);`49671 **DESCRIPTION**49672 **CX** The functionality described on this reference page is aligned with the ISO C standard. Any  
49673 conflict between the requirements described here and the ISO C standard is unintentional. This  
49674 volume of IEEE Std 1003.1-200x defers to the ISO C standard.49675 The `wcsncmp()` function shall compare not more than *n* wide-character codes (wide-character  
49676 codes that follow a null wide-character code are not compared) from the array pointed to by *ws1*  
49677 to the array pointed to by *ws2*.49678 The sign of a non-zero return value shall be determined by the sign of the difference between the  
49679 values of the first pair of wide-character codes that differ in the objects being compared. |49680 **RETURN VALUE**49681 Upon successful completion, `wcsncmp()` shall return an integer greater than, equal to, or less  
49682 than 0, if the possibly null-terminated array pointed to by *ws1* is greater than, equal to, or less  
49683 than the possibly null-terminated array pointed to by *ws2*, respectively.49684 **ERRORS**

49685 No errors are defined.

49686 **EXAMPLES**

49687 None.

49688 **APPLICATION USAGE**

49689 None.

49690 **RATIONALE**

49691 None.

49692 **FUTURE DIRECTIONS**

49693 None.

49694 **SEE ALSO**49695 `wscmp()`, the Base Definitions volume of IEEE Std 1003.1-200x, `<wchar.h>`49696 **CHANGE HISTORY**

49697 First released in Issue 4. Derived from the MSE working draft.

49698 **NAME**

49699       wcsncpy — copy part of a wide-character string

49700 **SYNOPSIS**

49701       #include &lt;wchar.h&gt;

49702       wchar\_t \*wcsncpy(wchar\_t \*restrict ws1, const wchar\_t \*restrict ws2,  
49703                       size\_t n);49704 **DESCRIPTION**49705 **CX**       The functionality described on this reference page is aligned with the ISO C standard. Any  
49706       conflict between the requirements described here and the ISO C standard is unintentional. This  
49707       volume of IEEE Std 1003.1-200x defers to the ISO C standard.49708       The *wcsncpy()* function shall copy not more than *n* wide-character codes (wide-character codes  
49709       that follow a null wide-character code are not copied) from the array pointed to by *ws2* to the  
49710       array pointed to by *ws1*. If copying takes place between objects that overlap, the behavior is  
49711       undefined.49712       If the array pointed to by *ws2* is a wide-character string that is shorter than *n* wide-character |  
49713       codes, null wide-character codes shall be appended to the copy in the array pointed to by *ws1*, |  
49714       until *n* wide-character codes in all are written.49715 **RETURN VALUE**49716       The *wcsncpy()* function shall return *ws1*; no return value is reserved to indicate an error.49717 **ERRORS**

49718       No errors are defined.

49719 **EXAMPLES**

49720       None.

49721 **APPLICATION USAGE**49722       If there is no null wide-character code in the first *n* wide-character codes of the array pointed to  
49723       by *ws2*, the result is not null-terminated.49724 **RATIONALE**

49725       None.

49726 **FUTURE DIRECTIONS**

49727       None.

49728 **SEE ALSO**49729       *wscpy()*, the Base Definitions volume of IEEE Std 1003.1-200x, <wchar.h>49730 **CHANGE HISTORY**

49731       First released in Issue 4. Derived from the MSE working draft.

49732 **Issue 6**49733       The *wcsncpy()* prototype is updated for alignment with the ISO/IEC 9899:1999 standard.

49734 **NAME**

49735 wcpbrk — scan wide-character string for a wide-character code

49736 **SYNOPSIS**

49737 #include &lt;wchar.h&gt;

49738 wchar\_t \*wcpbrk(const wchar\_t \*ws1, const wchar\_t \*ws2);

49739 **DESCRIPTION**

49740 cx The functionality described on this reference page is aligned with the ISO C standard. Any  
49741 conflict between the requirements described here and the ISO C standard is unintentional. This  
49742 volume of IEEE Std 1003.1-200x defers to the ISO C standard.

49743 The *wcpbrk()* function shall locate the first occurrence in the wide-character string pointed to by  
49744 *ws1* of any wide-character code from the wide-character string pointed to by *ws2*.

49745 **RETURN VALUE**

49746 Upon successful completion, *wcpbrk()* shall return a pointer to the wide-character code or a null  
49747 pointer if no wide-character code from *ws2* occurs in *ws1*.

49748 **ERRORS**

49749 No errors are defined.

49750 **EXAMPLES**

49751 None.

49752 **APPLICATION USAGE**

49753 None.

49754 **RATIONALE**

49755 None.

49756 **FUTURE DIRECTIONS**

49757 None.

49758 **SEE ALSO**49759 *wchr()*, *wchr()*, the Base Definitions volume of IEEE Std 1003.1-200x, <wchar.h>49760 **CHANGE HISTORY**

49761 First released in Issue 4. Derived from the MSE working draft.

49762 **NAME**

49763       wcsrchr — wide-character string scanning operation

49764 **SYNOPSIS**

49765       #include &lt;wchar.h&gt;

49766       wchar\_t \*wcsrchr(const wchar\_t \*ws, wchar\_t wc);

49767 **DESCRIPTION**49768 **CX**       The functionality described on this reference page is aligned with the ISO C standard. Any  
49769 conflict between the requirements described here and the ISO C standard is unintentional. This  
49770 volume of IEEE Std 1003.1-200x defers to the ISO C standard.49771       The *wcsrchr()* function shall locate the last occurrence of *wc* in the wide-character string pointed  
49772 to by *ws*. The application shall ensure that the value of *wc* is a character representable as a type  
49773 **wchar\_t** and a wide-character code corresponding to a valid character in the current locale. The  
49774 terminating null wide-character code shall be considered to be part of the wide-character string. |49775 **RETURN VALUE**49776       Upon successful completion, *wcsrchr()* shall return a pointer to the wide-character code or a null  
49777 pointer if *wc* does not occur in the wide-character string.49778 **ERRORS**

49779       No errors are defined.

49780 **EXAMPLES**

49781       None.

49782 **APPLICATION USAGE**

49783       None.

49784 **RATIONALE**

49785       None.

49786 **FUTURE DIRECTIONS**

49787       None.

49788 **SEE ALSO**49789       *wchr()*, the Base Definitions volume of IEEE Std 1003.1-200x, <**wchar.h**>49790 **CHANGE HISTORY**

49791       First released in Issue 4. Derived from the MSE working draft.

49792 **Issue 6**

49793       The DESCRIPTION is updated to avoid use of the term “must” for application requirements.

49794 **NAME**49795 `wcsrtombs` — convert a wide-character string to a character string (restartable)49796 **SYNOPSIS**49797 `#include <wchar.h>`49798 `size_t wcsrtombs(char *restrict dst, const wchar_t **restrict src,`  
49799 `size_t len, mbstate_t *restrict ps);`49800 **DESCRIPTION**49801 **CX** The functionality described on this reference page is aligned with the ISO C standard. Any  
49802 conflict between the requirements described here and the ISO C standard is unintentional. This  
49803 volume of IEEE Std 1003.1-200x defers to the ISO C standard.49804 The `wcsrtombs()` function shall convert a sequence of wide characters from the array indirectly  
49805 pointed to by `src` into a sequence of corresponding characters, beginning in the conversion state  
49806 described by the object pointed to by `ps`. If `dst` is not a null pointer, the converted characters  
49807 shall then be stored into the array pointed to by `dst`. Conversion continues up to and including a  
49808 terminating null wide character, which shall also be stored. Conversion shall stop earlier in the  
49809 following cases:

- 49810
- 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  
49811 pointed to by `dst` (and `dst` is not a null pointer)
- 49812

49813 Each conversion shall take place as if by a call to the `wcrtomb()` function. |49814 If `dst` is not a null pointer, the pointer object pointed to by `src` shall be assigned either a null |  
49815 pointer (if conversion stopped due to reaching a terminating null wide character) or the address |  
49816 just past the last wide character converted (if any). If conversion stopped due to reaching a |  
49817 terminating null wide character, the resulting state described shall be the initial conversion state. |49818 If `ps` is a null pointer, the `wcsrtombs()` function shall use its own internal `mbstate_t` object, which |  
49819 is initialized at program start-up to the initial conversion state. Otherwise, the `mbstate_t` object |  
49820 pointed to by `ps` shall be used to completely describe the current conversion state of the |  
49821 associated character sequence. The implementation shall behave as if no function defined in this |  
49822 volume of IEEE Std 1003.1-200x calls `wcsrtombs()`. |49823 **CX** If the application uses any of the `_POSIX_THREAD_SAFE_FUNCTIONS` or `_POSIX_THREADS`  
49824 functions, the application shall ensure that the `wcsrtombs()` function is called with a non-NULL  
49825 `ps` argument.49826 The behavior of this function shall be affected by the `LC_CTYPE` category of the current locale. |49827 **RETURN VALUE**49828 If conversion stops because a code is reached that does not correspond to a valid character, an  
49829 encoding error occurs. In this case, the `wcsrtombs()` function shall store the value of the macro  
49830 `[EILSEQ]` in `errno` and return `(size_t)-1`; the conversion state is undefined. Otherwise, it shall  
49831 return the number of bytes in the resulting character sequence, not including the terminating  
49832 null (if any).49833 **ERRORS**49834 The `wcsrtombs()` function may fail if:

- 49835
- CX**
- `[EINVAL]`
- `ps`
- points to an object that contains an invalid conversion state.
- 
- 49836
- `[EILSEQ]`
- A wide-character code does not correspond to a valid character.

49837 **EXAMPLES**

49838 None.

49839 **APPLICATION USAGE**

49840 None.

49841 **RATIONALE**

49842 None.

49843 **FUTURE DIRECTIONS**

49844 None.

49845 **SEE ALSO**49846 *mbsinit()*, *wcrtomb()*, the Base Definitions volume of IEEE Std 1003.1-200x, <**wchar.h**>49847 **CHANGE HISTORY**49848 First released in Issue 5. Included for alignment with ISO/IEC 9899:1990/Amendment 1:1995  
49849 (E).49850 **Issue 6**

49851 In the DESCRIPTION, a note on using this function in a threaded application is added.

49852 Extensions beyond the ISO C standard are now marked.

49853 The DESCRIPTION is updated to avoid use of the term “must” for application requirements.

49854 The *wcsrombs()* prototype is updated for alignment with the ISO/IEC 9899:1999 standard.



49855 **NAME**49856           wcssp<sub>n</sub> — get length of a wide substring49857 **SYNOPSIS**

49858           #include &lt;wchar.h&gt;

49859           size\_t wcssp<sub>n</sub>(const wchar\_t \*ws1, const wchar\_t \*ws2);49860 **DESCRIPTION**

49861 cx       The functionality described on this reference page is aligned with the ISO C standard. Any  
49862       conflict between the requirements described here and the ISO C standard is unintentional. This  
49863       volume of IEEE Std 1003.1-200x defers to the ISO C standard.

49864       The wcssp<sub>n</sub>() function shall compute the length (in wide characters) of the maximum initial |  
49865       segment of the wide-character string pointed to by ws1 which consists entirely of wide-character |  
49866       codes from the wide-character string pointed to by ws2.

49867 **RETURN VALUE**

49868       The wcssp<sub>n</sub>() function shall return the length of the initial substring of ws1; no return value is  
49869       reserved to indicate an error.

49870 **ERRORS**

49871       No errors are defined.

49872 **EXAMPLES**

49873       None.

49874 **APPLICATION USAGE**

49875       None.

49876 **RATIONALE**

49877       None.

49878 **FUTURE DIRECTIONS**

49879       None.

49880 **SEE ALSO**49881       wcscsp<sub>n</sub>(), the Base Definitions volume of IEEE Std 1003.1-200x, <wchar.h>49882 **CHANGE HISTORY**

49883       First released in Issue 4. Derived from the MSE working draft.

49884 **Issue 5**

49885       The RETURN VALUE section is updated to indicate that wcssp<sub>n</sub>() returns the length of ws1  
49886       rather than ws1 itself.

49887 **NAME**49888 `wcsstr` — find a wide-character substring49889 **SYNOPSIS**49890 `#include <wchar.h>`49891 `wchar_t *wcsstr(const wchar_t *restrict ws1, const wchar_t *restrict ws2);`49892 **DESCRIPTION**

49893 **CX** The functionality described on this reference page is aligned with the ISO C standard. Any  
49894 conflict between the requirements described here and the ISO C standard is unintentional. This  
49895 volume of IEEE Std 1003.1-200x defers to the ISO C standard.

49896 The `wcsstr()` function shall locate the first occurrence in the wide-character string pointed to by  
49897 `ws1` of the sequence of wide characters (excluding the terminating null wide character) in the  
49898 wide-character string pointed to by `ws2`.

49899 **RETURN VALUE**

49900 Upon successful completion, `wcsstr()` shall return a pointer to the located wide-character string,  
49901 or a null pointer if the wide-character string is not found.

49902 If `ws2` points to a wide-character string with zero length, the function shall return `ws1`.

49903 **ERRORS**

49904 No errors are defined.

49905 **EXAMPLES**

49906 None.

49907 **APPLICATION USAGE**

49908 None.

49909 **RATIONALE**

49910 None.

49911 **FUTURE DIRECTIONS**

49912 None.

49913 **SEE ALSO**49914 `wcschr()`, the Base Definitions volume of IEEE Std 1003.1-200x, `<wchar.h>`49915 **CHANGE HISTORY**

49916 First released in Issue 5. Included for alignment with ISO/IEC 9899:1990/Amendment 1:1995  
49917 (E).

49918 **Issue 6**49919 The `wcsstr()` prototype is updated for alignment with the ISO/IEC 9899:1999 standard.

## 49920 NAME

49921 `wcstod`, `wcstof`, `wcstold` — convert a wide-character string to a double-precision number

## 49922 SYNOPSIS

49923 `#include <wchar.h>`

49924 `double wcstod(const wchar_t *restrict nptr, wchar_t **restrict endptr);`

49925 `float wcstof(const wchar_t *restrict nptr, wchar_t **restrict endptr);`

49926 `long double wcstold(const wchar_t *restrict nptr,`

49927 `wchar_t **restrict endptr);`

## 49928 DESCRIPTION

49929 **CX** The functionality described on this reference page is aligned with the ISO C standard. Any  
 49930 conflict between the requirements described here and the ISO C standard is unintentional. This  
 49931 volume of IEEE Std 1003.1-200x defers to the ISO C standard.

49932 These functions shall convert the initial portion of the wide-character string pointed to by *nptr* to  
 49933 **double**, **float**, and **long double** representation, respectively. First, they shall decompose the  
 49934 input wide-character string into three parts:

- 49935 1. An initial, possibly empty, sequence of white-space wide-character codes (as specified by  
 49936 `iswspace()`)
- 49937 2. A subject sequence interpreted as a floating-point constant or representing infinity or NaN
- 49938 3. A final wide-character string of one or more unrecognized wide-character codes, including  
 49939 the terminating null wide-character code of the input wide-character string

49940 Then they shall attempt to convert the subject sequence to a floating-point number, and return  
 49941 the result.

49942 The expected form of the subject sequence is an optional plus or minus sign, then one of the  
 49943 following:

- 49944 • A non-empty sequence of decimal digits optionally containing a radix character, then an  
 49945 optional exponent part
- 49946 • A `0x` or `0X`, then a non-empty sequence of hexadecimal digits optionally containing a radix  
 49947 character, then an optional binary exponent part
- 49948 • One of `INF` or `INFINITY`, or any other wide string equivalent except for case
- 49949 • One of `NAN` or `NAN(n-wchar-sequenceopt)`, or any other wide string ignoring case in the NAN  
 49950 part, where:

49951 `n-wchar-sequence:`

49952 `digit`

49953 `nondigit`

49954 `n-wchar-sequence digit`

49955 `n-wchar-sequence nondigit`

49956 The subject sequence is defined as the longest initial subsequence of the input wide string,  
 49957 starting with the first non-white-space wide character, that is of the expected form. The subject  
 49958 sequence contains no wide characters if the input wide string is not of the expected form.

49959 If the subject sequence has the expected form for a floating-point number, the sequence of wide  
 49960 characters starting with the first digit or the radix character (whichever occurs first) shall be  
 49961 interpreted as a floating constant according to the rules of the C language, except that the radix  
 49962 character shall be used in place of a period, and that if neither an exponent part nor a radix  
 49963 character appears in a decimal floating-point number, or if a binary exponent part does not

- 49964 appear in a hexadecimal floating-point number, an exponent part of the appropriate type with  
 49965 value zero shall be assumed to follow the last digit in the string. If the subject sequence begins  
 49966 with a minus sign, the sequence shall be interpreted as negated. A wide-character sequence INF  
 49967 or INFINITY shall be interpreted as an infinity, if representable in the return type, else as if it  
 49968 were a floating constant that is too large for the range of the return type. A wide-character  
 49969 sequence NAN or NAN(*n-wchar-sequence<sub>opt</sub>*) shall be interpreted as a quiet NaN, if supported in  
 49970 the return type, else as if it were a subject sequence part that does not have the expected form;  
 49971 the meaning of the *n-wchar* sequences is implementation-defined. A pointer to the final wide  
 49972 string shall be stored in the object pointed to by *endptr*, provided that *endptr* is not a null pointer.
- 49973 If the subject sequence has the hexadecimal form and FLT\_RADIX is a power of 2, the  
 49974 conversion shall be rounded in an implementation-defined manner.
- 49975 CX The radix character shall be as defined in the program's locale (category *LC\_NUMERIC*). In the  
 49976 POSIX locale, or in a locale where the radix character is not defined, the radix character shall  
 49977 default to a period ( ' . ' ).
- 49978 CX In other than the C or POSIX locales, other implementation-defined subject sequences may be  
 49979 accepted.
- 49980 If the subject sequence is empty or does not have the expected form, no conversion shall be  
 49981 performed; the value of *nptr* shall be stored in the object pointed to by *endptr*, provided that  
 49982 *endptr* is not a null pointer.
- 49983 CX The *wcstod()* function shall not change the setting of *errno* if successful.
- 49984 Since 0 is returned on error and is also a valid return on success, an application wishing to check  
 49985 for error situations should set *errno* to 0, then call *wcstod()*, *wcstof()*, or *wcstold()*, then check  
 49986 *errno*.
- 49987 **RETURN VALUE**
- 49988 Upon successful completion, these functions shall return the converted value. If no conversion  
 49989 CX could be performed, 0 shall be returned and *errno* may be set to [EINVAL].
- 49990 If the correct value is outside the range of representable values, plus or minus HUGE\_VAL,  
 49991 HUGE\_VALF, or HUGE\_VALL shall be returned (according to the sign of the value), and *errno*  
 49992 shall be set to [ERANGE].
- 49993 If the correct value would cause underflow, a value whose magnitude is no greater than the  
 49994 smallest normalized positive number in the return type shall be returned and *errno* set to  
 49995 [ERANGE].
- 49996 **ERRORS**
- 49997 The *wcstod()* function shall fail if:
- 49998 [ERANGE] The value to be returned would cause overflow or underflow.
- 49999 The *wcstod()* function may fail if:
- 50000 CX [EINVAL] No conversion could be performed.

50001 **EXAMPLES**

50002 None.

50003 **APPLICATION USAGE**

50004 If the subject sequence has the hexadecimal form and FLT\_RADIX is not a power of 2, and the  
 50005 result is not exactly representable, the result should be one of the two numbers in the  
 50006 appropriate internal format that are adjacent to the hexadecimal floating source value, with the  
 50007 extra stipulation that the error should have a correct sign for the current rounding direction.

50008 If the subject sequence has the decimal form and at most DECIMAL\_DIG (defined in <float.h>)  
 50009 significant digits, the result should be correctly rounded. If the subject sequence *D* has the  
 50010 decimal form and more than DECIMAL\_DIG significant digits, consider the two bounding,  
 50011 adjacent decimal strings *L* and *U*, both having DECIMAL\_DIG significant digits, such that the  
 50012 values of *L*, *D*, and *U* satisfy " $L \leq D \leq U$ ". The result should be one of the (equal or  
 50013 adjacent) values that would be obtained by correctly rounding *L* and *U* according to the current  
 50014 rounding direction, with the extra stipulation that the error with respect to *D* should have a  
 50015 correct sign for the current rounding direction.

50016 **RATIONALE**

50017 None.

50018 **FUTURE DIRECTIONS**

50019 None.

50020 **SEE ALSO**

50021 *iswspace()*, *localeconv()*, *scanf()*, *setlocale()*, *wcstol()*, the Base Definitions volume of  
 50022 IEEE Std 1003.1-200x, <float.h>, <wchar.h>, the Base Definitions volume of  
 50023 IEEE Std 1003.1-200x, Chapter 7, Locale

50024 **CHANGE HISTORY**

50025 First released in Issue 4. Derived from the MSE working draft.

50026 **Issue 5**50027 The DESCRIPTION is updated to indicate that *errno* is not changed if the function is successful.50028 **Issue 6**

50029 Extensions beyond the ISO C standard are now marked.

50030 The following new requirements on POSIX implementations derive from alignment with the  
 50031 Single UNIX Specification:

50032 • In the RETURN VALUE and ERRORS sections, the [EINVAL] optional error condition is  
 50033 added if no conversion could be performed.

50034 The following changes are made for alignment with the ISO/IEC 9899:1999 standard:

50035 • The *wcstod()* prototype is updated.50036 • The *wcstof()* and *wcstold()* functions are added.

50037 • If the correct value for *wcstod()* would cause underflow, the return value changed from 0 (as  
 50038 specified in Issue 5) to the smallest normalized positive number.

50039 • The DESCRIPTION, RETURN VALUE, and APPLICATION USAGE sections are extensively  
 50040 updated.

50041 ISO/IEC 9899:1999 standard, Technical Corrigendum No. 1 is incorporated.

50042 **NAME**50043            **wcstof** — convert a wide-character string to a double-precision number50044 **SYNOPSIS**50045            `#include <wchar.h>`50046            `float wcstof(const wchar_t *restrict nptr, wchar_t **restrict endptr);`50047 **DESCRIPTION**50048            Refer to *wcstod()*.

50049 **NAME**50050 `wcstoimax`, `wcstoumax` — convert wide-character string to integer type50051 **SYNOPSIS**50052 `#include <stddef.h>`50053 `#include <inttypes.h>`50054 `intmax_t wcstoimax(const wchar_t *restrict nptr,`50055 `wchar_t **restrict endpnr, int base);`50056 `uintmax_t wcstoumax(const wchar_t *restrict nptr,`50057 `wchar_t **restrict endpnr, int base);`50058 **DESCRIPTION**50059 **CX** The functionality described on this reference page is aligned with the ISO C standard. Any  
50060 conflict between the requirements described here and the ISO C standard is unintentional. This  
50061 volume of IEEE Std 1003.1-200x defers to the ISO C standard.50062 These functions shall be equivalent to the `wcstol()`, `wcstoll()`, `wcstoul()`, and `wcstoull()` functions,  
50063 respectively, except that the initial portion of the wide string shall be converted to **intmax\_t** and  
50064 **uintmax\_t** representation, respectively.50065 **RETURN VALUE**

50066 These functions shall return the converted value, if any.

50067 If no conversion could be performed, zero shall be returned. If the correct value is outside the  
50068 range of representable values, {INTMAX\_MAX}, {INTMAX\_MIN}, or {UINTMAX\_MAX} shall  
50069 be returned (according to the return type and sign of the value, if any), and `errno` shall be set to  
50070 [ERANGE].50071 **ERRORS**

50072 These functions shall fail if:

50073 [EINVAL] The value of *base* is not supported.

50074 [ERANGE] The value to be returned is not representable.

50075 These functions may fail if:

50076 [EINVAL] No conversion could be performed.

50077 **EXAMPLES**

50078 None.

50079 **APPLICATION USAGE**

50080 None.

50081 **RATIONALE**

50082 None.

50083 **FUTURE DIRECTIONS**

50084 None.

50085 **SEE ALSO**50086 `wcstol()`, `wcstoul()`, the Base Definitions volume of IEEE Std 1003.1-200x, `<inttypes.h>`,50087 `<stddef.h>`50088 **CHANGE HISTORY**

50089 First released in Issue 6. Derived from the ISO/IEC 9899:1999 standard.

## 50090 NAME

50091 wcstok — split wide-character string into tokens

## 50092 SYNOPSIS

50093 #include &lt;wchar.h&gt;

50094 wchar\_t \*wcstok(wchar\_t \*restrict ws1, const wchar\_t \*restrict ws2,  
50095 wchar\_t \*\*restrict ptr);

## 50096 DESCRIPTION

50097 cx The functionality described on this reference page is aligned with the ISO C standard. Any  
50098 conflict between the requirements described here and the ISO C standard is unintentional. This  
50099 volume of IEEE Std 1003.1-200x defers to the ISO C standard.50100 A sequence of calls to *wcstok()* shall break the wide-character string pointed to by *ws1* into a |  
50101 sequence of tokens, each of which shall be delimited by a wide-character code from the wide- |  
50102 character string pointed to by *ws2*. The *ptr* argument points to a caller-provided **wchar\_t** pointer |  
50103 into which the *wcstok()* function shall store information necessary for it to continue scanning the |  
50104 same wide-character string. |50105 The first call in the sequence has *ws1* as its first argument, and is followed by calls with a null |  
50106 pointer as their first argument. The separator string pointed to by *ws2* may be different from call |  
50107 to call.50108 The first call in the sequence shall search the wide-character string pointed to by *ws1* for the first |  
50109 wide-character code that is *not* contained in the current separator string pointed to by *ws2*. If no |  
50110 such wide-character code is found, then there are no tokens in the wide-character string pointed |  
50111 to by *ws1* and *wcstok()* shall return a null pointer. If such a wide-character code is found, it shall |  
50112 be the start of the first token. |50113 The *wcstok()* function shall then search from there for a wide-character code that *is* contained in |  
50114 the current separator string. If no such wide-character code is found, the current token extends |  
50115 to the end of the wide-character string pointed to by *ws1*, and subsequent searches for a token |  
50116 shall return a null pointer. If such a wide-character code is found, it shall be overwritten by a |  
50117 null wide character, which terminates the current token. The *wcstok()* function shall save a |  
50118 pointer to the following wide-character code, from which the next search for a token shall start. |50119 Each subsequent call, with a null pointer as the value of the first argument, shall start searching |  
50120 from the saved pointer and behave as described above. |50121 The implementation shall behave as if no function calls *wcstok()*.

## 50122 RETURN VALUE

50123 Upon successful completion, the *wcstok()* function shall return a pointer to the first wide- |  
50124 character code of a token. Otherwise, if there is no token, *wcstok()* shall return a null pointer.

## 50125 ERRORS

50126 No errors are defined.



50127 **EXAMPLES**

50128 None.

50129 **APPLICATION USAGE**

50130 None.

50131 **RATIONALE**

50132 None.

50133 **FUTURE DIRECTIONS**

50134 None.

50135 **SEE ALSO**

50136 The Base Definitions volume of IEEE Std 1003.1-200x, &lt;wchar.h&gt;

50137 **CHANGE HISTORY**

50138 First released in Issue 4.

50139 **Issue 5**50140 Aligned with ISO/IEC 9899:1990/Amendment 1:1995 (E). Specifically, a third argument is  
50141 added to the definition of this function in the SYNOPSIS.50142 **Issue 6**50143 The *wcstok()* prototype is updated for alignment with the ISO/IEC 9899:1999 standard.

## 50144 NAME

50145 wcstol, wcstoll — convert a wide-character string to a long integer

## 50146 SYNOPSIS

50147 #include <wchar.h>

50148 long wcstol(const wchar\_t \*restrict *nptr*, wchar\_t \*\*restrict *endptr*,  
50149 int *base*);

50150 long long wcstoll(const wchar\_t \*restrict *nptr*,  
50151 wchar\_t \*\*restrict *endptr*, int *base*);

## 50152 DESCRIPTION

50153 cx The functionality described on this reference page is aligned with the ISO C standard. Any  
50154 conflict between the requirements described here and the ISO C standard is unintentional. This  
50155 volume of IEEE Std 1003.1-200x defers to the ISO C standard.

50156 These functions shall convert the initial portion of the wide-character string pointed to by *nptr* to  
50157 **long**, **long long**, **unsigned long**, and **unsigned long long** representation, respectively. First, they  
50158 shall decompose the input string into three parts:

- 50159 1. An initial, possibly empty, sequence of white-space wide-character codes (as specified by  
50160 *iswspace()*)
- 50161 2. A subject sequence interpreted as an integer represented in some radix determined by the  
50162 value of *base*
- 50163 3. A final wide-character string of one or more unrecognized wide-character codes, including  
50164 the terminating null wide-character code of the input wide-character string

50165 Then they shall attempt to convert the subject sequence to an integer, and return the result.

50166 If *base* is 0, the expected form of the subject sequence is that of a decimal constant, octal constant,  
50167 or hexadecimal constant, any of which may be preceded by a '+' or '-' sign. A decimal  
50168 constant begins with a non-zero digit, and consists of a sequence of decimal digits. An octal  
50169 constant consists of the prefix '0' optionally followed by a sequence of the digits '0' to '7'  
50170 only. A hexadecimal constant consists of the prefix 0x or 0X followed by a sequence of the  
50171 decimal digits and letters 'a' (or 'A') to 'f' (or 'F') with values 10 to 15 respectively.

50172 If the value of *base* is between 2 and 36, the expected form of the subject sequence is a sequence  
50173 of letters and digits representing an integer with the radix specified by *base*, optionally preceded  
50174 by a '+' or '-' sign, but not including an integer suffix. The letters from 'a' (or 'A') to 'z'  
50175 (or 'Z') inclusive are ascribed the values 10 to 35; only letters whose ascribed values are less  
50176 than that of *base* shall be permitted. If the value of *base* is 16, the wide-character code  
50177 representations of 0x or 0X may optionally precede the sequence of letters and digits, following  
50178 the sign if present.

50179 The subject sequence is defined as the longest initial subsequence of the input wide-character  
50180 string, starting with the first non-white-space wide-character code that is of the expected form.  
50181 The subject sequence contains no wide-character codes if the input wide-character string is  
50182 empty or consists entirely of white-space wide-character code, or if the first non-white-space  
50183 wide-character code is other than a sign or a permissible letter or digit.

50184 If the subject sequence has the expected form and *base* is 0, the sequence of wide-character codes  
50185 starting with the first digit shall be interpreted as an integer constant. If the subject sequence has  
50186 the expected form and the value of *base* is between 2 and 36, it shall be used as the base for  
50187 conversion, ascribing to each letter its value as given above. If the subject sequence begins with a  
50188 minus sign, the value resulting from the conversion shall be negated. A pointer to the final  
50189 wide-character string shall be stored in the object pointed to by *endptr*, provided that *endptr* is

- 50190 not a null pointer.
- 50191 CX In other than the C or POSIX locales, other implementation-defined subject sequences may be  
50192 accepted.
- 50193 If the subject sequence is empty or does not have the expected form, no conversion shall be |  
50194 performed; the value of *nptr* shall be stored in the object pointed to by *endptr*, provided that |  
50195 *endptr* is not a null pointer.
- 50196 CX These functions shall not change the setting of *errno* if successful.
- 50197 Since 0, {LONG\_MIN} or {LLONG\_MIN} and {LONG\_MAX} or {LLONG\_MAX} are returned on |  
50198 error and are also valid returns on success, an application wishing to check for error situations  
50199 should set *errno* to 0, then call *wcstol()* or *wcstoll()*, then check *errno*.
- 50200 **RETURN VALUE**
- 50201 Upon successful completion, these functions shall return the converted value, if any. If no  
50202 CX conversion could be performed, 0 shall be returned and *errno* may be set to indicate the error. If  
50203 the correct value is outside the range of representable values, {LONG\_MIN}, {LONG\_MAX},  
50204 {LLONG\_MIN}, or {LLONG\_MAX} shall be returned (according to the sign of the value), and  
50205 *errno* set to [ERANGE].
- 50206 **ERRORS**
- 50207 These functions shall fail if:
- 50208 CX [EINVAL] The value of *base* is not supported.
- 50209 [ERANGE] The value to be returned is not representable.
- 50210 These functions may fail if:
- 50211 CX [EINVAL] No conversion could be performed.
- 50212 **EXAMPLES**
- 50213 None.
- 50214 **APPLICATION USAGE**
- 50215 None.
- 50216 **RATIONALE**
- 50217 None.
- 50218 **FUTURE DIRECTIONS**
- 50219 None.
- 50220 **SEE ALSO**
- 50221 *iswalpha()*, *scanf()*, *wcstod()*, the Base Definitions volume of IEEE Std 1003.1-200x, <**wchar.h**>
- 50222 **CHANGE HISTORY**
- 50223 First released in Issue 4. Derived from the MSE working draft.
- 50224 **Issue 5**
- 50225 The DESCRIPTION is updated to indicate that *errno* is not changed if the function is successful.
- 50226 **Issue 6**
- 50227 Extensions beyond the ISO C standard are now marked.
- 50228 The following new requirements on POSIX implementations derive from alignment with the  
50229 Single UNIX Specification:
- 50230 • In the RETURN VALUE and ERRORS sections, the [EINVAL] optional error condition is  
50231 added if no conversion could be performed.

50232 The following changes are made for alignment with the ISO/IEC 9899:1999 standard:

50233 • The *wcstol()* prototype is updated.

50234 • The *wcstoll()* function is added.

50235 **NAME**

50236           wcstold — convert a wide-character string to a double-precision number

50237 **SYNOPSIS**

50238           #include <wchar.h>

50239           long double wcstold(const wchar\_t \*restrict *nptr*,  
50240                               wchar\_t \*\*restrict *endptr*);

50241 **DESCRIPTION**

50242           Refer to *wcstod*().

50243 **NAME**

50244           wcstoll — convert a wide-character string to a long integer

50245 **SYNOPSIS**

50246           #include <wchar.h>

50247           long long wcstoll(const wchar\_t \*restrict *nptr*,

50248                    wchar\_t \*\*restrict *endptr*, int *base*);

50249 **DESCRIPTION**

50250           Refer to *wcstol*().

50251 **NAME**50252 `wcstombs` — convert a wide-character string to a character string50253 **SYNOPSIS**50254 `#include <stdlib.h>`50255 `size_t wcstombs(char *restrict s, const wchar_t *restrict pwcs,`  
50256 `size_t n);`50257 **DESCRIPTION**50258 **CX** The functionality described on this reference page is aligned with the ISO C standard. Any  
50259 conflict between the requirements described here and the ISO C standard is unintentional. This  
50260 volume of IEEE Std 1003.1-200x defers to the ISO C standard.50261 The `wcstombs()` function shall convert the sequence of wide-character codes that are in the array  
50262 pointed to by `pwcs` into a sequence of characters that begins in the initial shift state and store  
50263 these characters into the array pointed to by `s`, stopping if a character would exceed the limit of `n`  
50264 total bytes or if a null byte is stored. Each wide-character code shall be converted as if by a call to  
50265 `wctomb()`, except that the shift state of `wctomb()` shall not be affected.50266 The behavior of this function shall be affected by the `LC_CTYPE` category of the current locale.50267 No more than `n` bytes shall be modified in the array pointed to by `s`. If copying takes place  
50268 **CX** between objects that overlap, the behavior is undefined. If `s` is a null pointer, `wcstombs()` shall  
50269 return the length required to convert the entire array regardless of the value of `n`, but no values  
50270 are stored.50271 The `wcstombs()` function need not be reentrant. A function that is not required to be reentrant is  
50272 not required to be thread-safe.50273 **RETURN VALUE**50274 If a wide-character code is encountered that does not correspond to a valid character (of one or  
50275 more bytes each), `wcstombs()` shall return `(size_t)-1`. Otherwise, `wcstombs()` shall return the  
50276 number of bytes stored in the character array, not including any terminating null byte. The array  
50277 shall not be null-terminated if the value returned is `n`.50278 **ERRORS**50279 The `wcstombs()` function may fail if:50280 **CX** `[EILSEQ]` A wide-character code does not correspond to a valid character.50281 **EXAMPLES**

50282 None.

50283 **APPLICATION USAGE**

50284 None.

50285 **RATIONALE**

50286 None.

50287 **FUTURE DIRECTIONS**

50288 None.

50289 **SEE ALSO**50290 `mblen()`, `mbtowc()`, `mbstowcs()`, `wctomb()`, the Base Definitions volume of IEEE Std 1003.1-200x,  
50291 `<stdlib.h>`

50292 **CHANGE HISTORY**

50293 First released in Issue 4. Derived from the ISO C standard.

50294 **Issue 6**

50295 The following new requirements on POSIX implementations derive from alignment with the  
50296 Single UNIX Specification:

- 50297 • The DESCRIPTION states the effect of when *s* is a null pointer.
- 50298 • The [EILSEQ] error condition is added.

50299 The *wcstombs()* prototype is updated for alignment with the ISO/IEC 9899:1999 standard.



## 50300 NAME

50301 wcstoul, wcstoull — convert a wide-character string to an unsigned long

## 50302 SYNOPSIS

50303 #include <wchar.h>

50304 unsigned long wcstoul(const wchar\_t \*restrict *nptr*,  
50305 wchar\_t \*\*restrict *endptr*, int *base*);

50306 unsigned long long wcstoull(const wchar\_t \*restrict *nptr*,  
50307 wchar\_t \*\*restrict *endptr*, int *base*);

## 50308 DESCRIPTION

50309 cx The functionality described on this reference page is aligned with the ISO C standard. Any  
50310 conflict between the requirements described here and the ISO C standard is unintentional. This  
50311 volume of IEEE Std 1003.1-200x defers to the ISO C standard.

50312 The *wcstoul()* and *wcstoull()* functions shall convert the initial portion of the wide-character  
50313 string pointed to by *nptr* to **unsigned long** and **unsigned long long** representation, respectively. |  
50314 First, they shall decompose the input wide-character string into three parts: |

- 50315 1. An initial, possibly empty, sequence of white-space wide-character codes (as specified by  
50316 *iswspace()*)
- 50317 2. A subject sequence interpreted as an integer represented in some radix determined by the  
50318 value of *base*
- 50319 3. A final wide-character string of one or more unrecognized wide-character codes, including  
50320 the terminating null wide-character code of the input wide-character string

50321 Then they shall attempt to convert the subject sequence to an unsigned integer, and return the |  
50322 result. |

50323 If *base* is 0, the expected form of the subject sequence is that of a decimal constant, octal constant,  
50324 or hexadecimal constant, any of which may be preceded by a '+' or '-' sign. A decimal  
50325 constant begins with a non-zero digit, and consists of a sequence of decimal digits. An octal  
50326 constant consists of the prefix '0' optionally followed by a sequence of the digits '0' to '7'  
50327 only. A hexadecimal constant consists of the prefix 0x or 0X followed by a sequence of the  
50328 decimal digits and letters 'a' (or 'A') to 'f' (or 'F') with values 10 to 15 respectively.

50329 If the value of *base* is between 2 and 36, the expected form of the subject sequence is a sequence  
50330 of letters and digits representing an integer with the radix specified by *base*, optionally preceded  
50331 by a '+' or '-' sign, but not including an integer suffix. The letters from 'a' (or 'A') to 'z'  
50332 (or 'Z') inclusive are ascribed the values 10 to 35; only letters whose ascribed values are less  
50333 than that of *base* shall be permitted. If the value of *base* is 16, the wide-character codes 0x or 0X |  
50334 may optionally precede the sequence of letters and digits, following the sign if present.

50335 The subject sequence is defined as the longest initial subsequence of the input wide-character  
50336 string, starting with the first wide-character code that is not white space and is of the expected  
50337 form. The subject sequence contains no wide-character codes if the input wide-character string is  
50338 empty or consists entirely of white-space wide-character codes, or if the first wide-character  
50339 code that is not white space is other than a sign or a permissible letter or digit.

50340 If the subject sequence has the expected form and *base* is 0, the sequence of wide-character codes  
50341 starting with the first digit shall be interpreted as an integer constant. If the subject sequence has |  
50342 the expected form and the value of *base* is between 2 and 36, it shall be used as the base for |  
50343 conversion, ascribing to each letter its value as given above. If the subject sequence begins with a |  
50344 minus sign, the value resulting from the conversion shall be negated. A pointer to the final |  
50345 wide-character string shall be stored in the object pointed to by *endptr*, provided that *endptr* is |

- 50346 not a null pointer.
- 50347 CX In other than the C or POSIX locales, other implementation-defined subject sequences may be  
50348 accepted.
- 50349 If the subject sequence is empty or does not have the expected form, no conversion shall be |  
50350 performed; the value of *nptr* shall be stored in the object pointed to by *endptr*, provided that |  
50351 *endptr* is not a null pointer.
- 50352 CX The *wcstoul()* function shall not change the setting of *errno* if successful.
- 50353 Since 0, {ULONG\_MAX}, and {ULLONG\_MAX} are returned on error and 0 is also a valid return |  
50354 on success, an application wishing to check for error situations should set *errno* to 0, then call |  
50355 *wcstoul()* or *wcstoull()*, then check *errno*.
- 50356 **RETURN VALUE**
- 50357 Upon successful completion, the *wcstoul()* and *wcstoull()* functions shall return the converted  
50358 CX value, if any. If no conversion could be performed, 0 shall be returned and *errno* may be set to  
50359 indicate the error. If the correct value is outside the range of representable values,  
50360 {ULONG\_MAX} or {ULLONG\_MAX} respectively shall be returned and *errno* set to [ERANGE].
- 50361 **ERRORS**
- 50362 These functions shall fail if:
- 50363 CX [EINVAL] The value of *base* is not supported.
- 50364 [ERANGE] The value to be returned is not representable.
- 50365 These functions may fail if:
- 50366 CX [EINVAL] No conversion could be performed.
- 50367 **EXAMPLES**
- 50368 None.
- 50369 **APPLICATION USAGE**
- 50370 None.
- 50371 **RATIONALE**
- 50372 None.
- 50373 **FUTURE DIRECTIONS**
- 50374 None.
- 50375 **SEE ALSO**
- 50376 *iswalph()*, *scanf()*, *wcstod()*, *wcstol()*, the Base Definitions volume of IEEE Std 1003.1-200x,  
50377 <wchar.h>
- 50378 **CHANGE HISTORY**
- 50379 First released in Issue 4. Derived from the MSE working draft.
- 50380 **Issue 5**
- 50381 The DESCRIPTION is updated to indicate that *errno* is not changed if the function is successful.
- 50382 **Issue 6**
- 50383 Extensions beyond the ISO C standard are now marked.
- 50384 The following new requirements on POSIX implementations derive from alignment with the  
50385 Single UNIX Specification:
- 50386
- The [EINVAL] error condition is added for when the value of *base* is not supported.

50387            In the RETURN VALUE and ERRORS sections, the [EINVAL] optional error condition is  
50388            added if no conversion could be performed.

50389            The following changes are made for alignment with the ISO/IEC 9899:1999 standard:

- 50390            • The *wcstoul()* prototype is updated.
- 50391            • The *wcstoull()* function is added.

50392 **NAME**

50393           wcstoull — convert a wide-character string to an unsigned long

50394 **SYNOPSIS**

50395           #include &lt;wchar.h&gt;

50396           unsigned long long wcstoull(const wchar\_t \*restrict *nptr*,50397                    wchar\_t \*\*restrict *endptr*, int *base*);50398 **DESCRIPTION**50399           Refer to *wcstoul()*.

50400 **NAME**

50401           wcstoumax — convert wide-character string to integer type

50402 **SYNOPSIS**

50403           #include <stddef.h>

50404           #include <inttypes.h>

50405           uintmax\_t wcstoumax(const wchar\_t \*restrict *nptr*,

50406                            wchar\_t \*\*restrict *endptr*, int *base*);

50407 **DESCRIPTION**

50408           Refer to *wcstoimax()*.

50409 **NAME**50410 wcswcs — find a wide substring (**LEGACY**)50411 **SYNOPSIS**50412 XSI `#include <wchar.h>`50413 `wchar_t *wcswcs(const wchar_t *ws1, const wchar_t *ws2);`

50414

50415 **DESCRIPTION**

50416 The `wcswcs()` function shall locate the first occurrence in the wide-character string pointed to by  
50417 `ws1` of the sequence of wide-character codes (excluding the terminating null wide-character  
50418 code) in the wide-character string pointed to by `ws2`.

50419 **RETURN VALUE**

50420 Upon successful completion, `wcswcs()` shall return a pointer to the located wide-character string  
50421 or a null pointer if the wide-character string is not found.

50422 If `ws2` points to a wide-character string with zero length, the function shall return `ws1`.

50423 **ERRORS**

50424 No errors are defined.

50425 **EXAMPLES**

50426 None.

50427 **APPLICATION USAGE**

50428 This function was not included in the final ISO/IEC 9899:1990/Amendment 1:1995 (E).  
50429 Application developers are strongly encouraged to use the `wcsstr()` function instead.

50430 **RATIONALE**

50431 None.

50432 **FUTURE DIRECTIONS**

50433 This function may be withdrawn in a future version.

50434 **SEE ALSO**50435 `wcschr()`, `wcsstr()`, the Base Definitions volume of IEEE Std 1003.1-200x, `<wchar.h>`50436 **CHANGE HISTORY**

50437 First released in Issue 4. Derived from the MSE working draft.

50438 **Issue 5**

50439 Marked EX.

50440 **Issue 6**

50441 This function is marked LEGACY.

50442 **NAME**50443        `wcswidth` — number of column positions of a wide-character string50444 **SYNOPSIS**50445 `XSI`        `#include <wchar.h>`50446        `int wcswidth(const wchar_t *pwcs, size_t n);`

50447

50448 **DESCRIPTION**

50449        The `wcswidth()` function shall determine the number of column positions required for *n* wide-character codes (or fewer than *n* wide-character codes if a null wide-character code is encountered before *n* wide-character codes are exhausted) in the string pointed to by *pwcs*.

50452 **RETURN VALUE**

50453        The `wcswidth()` function either shall return 0 (if *pwcs* points to a null wide-character code), or return the number of column positions to be occupied by the wide-character string pointed to by *pwcs*, or return -1 (if any of the first *n* wide-character codes in the wide-character string pointed to by *pwcs* is not a printable wide-character code).

50457 **ERRORS**

50458        No errors are defined.

50459 **EXAMPLES**

50460        None.

50461 **APPLICATION USAGE**

50462        This function was removed from the final ISO/IEC 9899:1990/Amendment 1:1995 (E), and the return value for a non-printable wide character is not specified.

50464 **RATIONALE**

50465        None.

50466 **FUTURE DIRECTIONS**

50467        None.

50468 **SEE ALSO**

50469        `wcwidth()`, the Base Definitions volume of IEEE Std 1003.1-200x, `<wchar.h>`, the Base Definitions volume of IEEE Std 1003.1-200x, Section 3.103, Column Position

50471 **CHANGE HISTORY**

50472        First released in Issue 4. Derived from the MSE working draft.

50473 **Issue 6**

50474        The Open Group Corrigendum U021/11 is applied. The function is marked as an extension.

50475 **NAME**50476 `wcsxfrm` — wide-character string transformation50477 **SYNOPSIS**50478 `#include <wchar.h>`50479 `size_t wcsxfrm(wchar_t *restrict ws1, const wchar_t *restrict ws2,`  
50480 `size_t n);`50481 **DESCRIPTION**50482 **CX** The functionality described on this reference page is aligned with the ISO C standard. Any  
50483 conflict between the requirements described here and the ISO C standard is unintentional. This  
50484 volume of IEEE Std 1003.1-200x defers to the ISO C standard.50485 The `wcsxfrm()` function shall transform the wide-character string pointed to by `ws2` and place the  
50486 resulting wide-character string into the array pointed to by `ws1`. The transformation shall be  
50487 such that if `wscmp()` is applied to two transformed wide strings, it shall return a value greater  
50488 than, equal to, or less than 0, corresponding to the result of `wscoll()` applied to the same two  
50489 original wide-character strings. No more than `n` wide-character codes shall be placed into the  
50490 resulting array pointed to by `ws1`, including the terminating null wide-character code. If `n` is 0,  
50491 `ws1` is permitted to be a null pointer. If copying takes place between objects that overlap, the  
50492 behavior is undefined.50493 **CX** The `wcsxfrm()` function shall not change the setting of `errno` if successful.50494 Since no return value is reserved to indicate an error, an application wishing to check for error  
50495 situations should set `errno` to 0, then call `wcsxfrm()`, then check `errno`.50496 **RETURN VALUE**50497 The `wcsxfrm()` function shall return the length of the transformed wide-character string (not  
50498 including the terminating null wide-character code). If the value returned is `n` or more, the  
50499 contents of the array pointed to by `ws1` are unspecified.50500 **CX** On error, the `wcsxfrm()` function may set `errno`, but no return value is reserved to indicate an  
50501 error.50502 **ERRORS**50503 The `wcsxfrm()` function may fail if:50504 **CX** [EINVAL] The wide-character string pointed to by `ws2` contains wide-character codes  
50505 outside the domain of the collating sequence.50506 **EXAMPLES**

50507 None.

50508 **APPLICATION USAGE**50509 The transformation function is such that two transformed wide-character strings can be ordered  
50510 by `wscmp()` as appropriate to collating sequence information in the program's locale (category  
50511 `LC_COLLATE`).50512 The fact that when `n` is 0 `ws1` is permitted to be a null pointer is useful to determine the size of  
50513 the `ws1` array prior to making the transformation.50514 **RATIONALE**

50515 None.



50516 **FUTURE DIRECTIONS**

50517 None.

50518 **SEE ALSO**50519 `wscmp()`, `wscoll()`, the Base Definitions volume of IEEE Std 1003.1-200x, <**wchar.h**>50520 **CHANGE HISTORY**

50521 First released in Issue 4. Derived from the MSE working draft.

50522 **Issue 5**

50523 Moved from ENHANCED I18N to BASE and the [ENOSYS] error is removed.

50524 The DESCRIPTION is updated to indicate that *errno* is not changed if the function is successful.50525 **Issue 6**

50526 In previous versions, this function was required to return -1 on error.

50527 Extensions beyond the ISO C standard are now marked.

50528 The following new requirements on POSIX implementations derive from alignment with the  
50529 Single UNIX Specification:

- 50530
- In the RETURN VALUE and ERRORS sections, the [EINVAL] optional error condition is
- 50531 added if no conversion could be performed.

50532 The `wcsxfrm()` prototype is updated for alignment with the ISO/IEC 9899:1999 standard.

50533 **NAME**

50534 wctob — wide-character to single-byte conversion

50535 **SYNOPSIS**

50536 #include <stdio.h>

50537 #include <wchar.h>

50538 int wctob(wint\_t c);

50539 **DESCRIPTION**

50540 cx The functionality described on this reference page is aligned with the ISO C standard. Any  
50541 conflict between the requirements described here and the ISO C standard is unintentional. This  
50542 volume of IEEE Std 1003.1-200x defers to the ISO C standard.

50543 The *wctob()* function shall determine whether *c* corresponds to a member of the extended  
50544 character set whose character representation is a single byte when in the initial shift state.

50545 The behavior of this function shall be affected by the *LC\_CTYPE* category of the current locale.

50546 **RETURN VALUE**

50547 The *wctob()* function shall return EOF if *c* does not correspond to a character with length one in  
50548 the initial shift state. Otherwise, it shall return the single-byte representation of that character as  
50549 an **unsigned char** converted to **int**.

50550 **ERRORS**

50551 No errors are defined.

50552 **EXAMPLES**

50553 None.

50554 **APPLICATION USAGE**

50555 None.

50556 **RATIONALE**

50557 None.

50558 **FUTURE DIRECTIONS**

50559 None.

50560 **SEE ALSO**

50561 *btowc()*, the Base Definitions volume of IEEE Std 1003.1-200x, <wchar.h>

50562 **CHANGE HISTORY**

50563 First released in Issue 5. Included for alignment with ISO/IEC 9899:1990/Amendment 1:1995  
50564 (E).

50565 **NAME**

50566 wctomb — convert a wide-character code to a character

50567 **SYNOPSIS**

50568 #include &lt;stdlib.h&gt;

50569 int wctomb(char \*s, wchar\_t wchar);

50570 **DESCRIPTION**

50571 cx The functionality described on this reference page is aligned with the ISO C standard. Any  
50572 conflict between the requirements described here and the ISO C standard is unintentional. This  
50573 volume of IEEE Std 1003.1-200x defers to the ISO C standard.

50574 The *wctomb()* function shall determine the number of bytes needed to represent the character  
50575 corresponding to the wide-character code whose value is *wchar* (including any change in the  
50576 shift state). It shall store the character representation (possibly multiple bytes and any special  
50577 bytes to change shift state) in the array object pointed to by *s* (if *s* is not a null pointer). At most  
50578 {MB\_CUR\_MAX} bytes shall be stored. If *wchar* is 0, a null byte shall be stored, preceded by any  
50579 shift sequence needed to restore the initial shift state, and *wctomb()* shall be left in the initial shift  
50580 state.

50581 cx The behavior of this function is affected by the *LC\_CTYPE* category of the current locale. For a  
50582 state-dependent encoding, this function shall be placed into its initial state by a call for which its  
50583 character pointer argument, *s*, is a null pointer. Subsequent calls with *s* as other than a null  
50584 pointer shall cause the internal state of the function to be altered as necessary. A call with *s* as a  
50585 null pointer shall cause this function to return a non-zero value if encodings have state  
50586 dependency, and 0 otherwise. Changing the *LC\_CTYPE* category causes the shift state of this  
50587 function to be unspecified.

50588 The *wctomb()* function need not be reentrant. A function that is not required to be reentrant is  
50589 not required to be thread-safe.

50590 The implementation shall behave as if no function defined in this volume of  
50591 IEEE Std 1003.1-200x calls *wctomb()*.

50592 **RETURN VALUE**

50593 If *s* is a null pointer, *wctomb()* shall return a non-zero or 0 value, if character encodings,  
50594 respectively, do or do not have state-dependent encodings. If *s* is not a null pointer, *wctomb()*  
50595 shall return -1 if the value of *wchar* does not correspond to a valid character, or return the  
50596 number of bytes that constitute the character corresponding to the value of *wchar*.

50597 In no case shall the value returned be greater than the value of the {MB\_CUR\_MAX} macro.

50598 **ERRORS**

50599 No errors are defined.

50600 **EXAMPLES**

50601 None.

50602 **APPLICATION USAGE**

50603 None.

50604 **RATIONALE**

50605 None.

50606 **FUTURE DIRECTIONS**

50607 None.

50608 **SEE ALSO**

50609            *mblen()*, *mbtowc()*, *mbstowcs()*, *wcstombs()*, the Base Definitions volume of IEEE Std 1003.1-200x,  
50610            <**stdlib.h**>

50611 **CHANGE HISTORY**

50612            First released in Issue 4. Derived from the ANSI C standard.

50613 **Issue 6**

50614            Extensions beyond the ISO C standard are now marked.

50615            In the DESCRIPTION, a note about reentrancy and thread-safety is added.

50616 **NAME**

50617 wctrans — define character mapping

50618 **SYNOPSIS**

50619 #include &lt;wctype.h&gt;

50620 wctrans\_t wctrans(const char \*charclass);

50621 **DESCRIPTION**

50622 cx The functionality described on this reference page is aligned with the ISO C standard. Any  
 50623 conflict between the requirements described here and the ISO C standard is unintentional. This  
 50624 volume of IEEE Std 1003.1-200x defers to the ISO C standard.

50625 The *wctrans()* function is defined for valid character mapping names identified in the current  
 50626 locale. The *charclass* is a string identifying a generic character mapping name for which codeset-  
 50627 specific information is required. The following character mapping names are defined in all  
 50628 locales: **tolower** and **toupper**.

50629 The function shall return a value of type **wctrans\_t**, which can be used as the second argument  
 50630 to subsequent calls of *towctrans()*. The *wctrans()* function shall determine values of **wctrans\_t**  
 50631 according to the rules of the coded character set defined by character mapping information in  
 50632 the program's locale (category *LC\_CTYPE*). The values returned by *wctrans()* shall be valid until  
 50633 a call to *setlocale()* that modifies the category *LC\_CTYPE*.

50634 **RETURN VALUE**

50635 cx The *wctrans()* function shall return 0 and may set *errno* to indicate the error if the given  
 50636 character mapping name is not valid for the current locale (category *LC\_CTYPE*); otherwise, it  
 50637 shall return a non-zero object of type **wctrans\_t** that can be used in calls to *towctrans()*.

50638 **ERRORS**50639 The *wctrans()* function may fail if:

50640 cx [EINVAL] The character mapping name pointed to by *charclass* is not valid in the current  
 50641 locale.

50642 **EXAMPLES**

50643 None.

50644 **APPLICATION USAGE**

50645 None.

50646 **RATIONALE**

50647 None.

50648 **FUTURE DIRECTIONS**

50649 None.

50650 **SEE ALSO**50651 *towctrans()*, the Base Definitions volume of IEEE Std 1003.1-200x, <wctype.h>50652 **CHANGE HISTORY**

50653 First released in Issue 5. Derived from ISO/IEC 9899:1990/Amendment 1:1995 (E).

50654 **NAME**

50655 wctype — define character class

50656 **SYNOPSIS**

50657 #include &lt;wctype.h&gt;

50658 wctype\_t wctype(const char \*property);

50659 **DESCRIPTION**

50660 cx The functionality described on this reference page is aligned with the ISO C standard. Any  
 50661 conflict between the requirements described here and the ISO C standard is unintentional. This  
 50662 volume of IEEE Std 1003.1-200x defers to the ISO C standard.

50663 The *wctype()* function is defined for valid character class names as defined in the current locale.  
 50664 The *property* argument is a string identifying a generic character class for which codeset-specific  
 50665 type information is required. The following character class names shall be defined in all locales:

50666	<b>alnum</b>	<b>digit</b>	<b>punct</b>
50667	<b>alpha</b>	<b>graph</b>	<b>space</b>
50668	<b>blank</b>	<b>lower</b>	<b>upper</b>
50669	<b>cntrl</b>	<b>print</b>	<b>xdigit</b>

50670 Additional character class names defined in the locale definition file (category *LC\_CTYPE*) can  
 50671 also be specified.

50672 The function shall return a value of type **wctype\_t**, which can be used as the second argument to  
 50673 subsequent calls of *iswctype()*. The *wctype()* function shall determine values of **wctype\_t**  
 50674 according to the rules of the coded character set defined by character type information in the  
 50675 program's locale (category *LC\_CTYPE*). The values returned by *wctype()* shall be valid until a  
 50676 call to *setlocale()* that modifies the category *LC\_CTYPE*.

50677 **RETURN VALUE**

50678 The *wctype()* function shall return 0 if the given character class name is not valid for the current  
 50679 locale (category *LC\_CTYPE*); otherwise, it shall return an object of type **wctype\_t** that can be  
 50680 used in calls to *iswctype()*.

50681 **ERRORS**

50682 No errors are defined.

50683 **EXAMPLES**

50684 None.

50685 **APPLICATION USAGE**

50686 None.

50687 **RATIONALE**

50688 None.

50689 **FUTURE DIRECTIONS**

50690 None.

50691 **SEE ALSO**50692 *iswctype()*, the Base Definitions volume of IEEE Std 1003.1-200x, <wctype.h>, <wchar.h>50693 **CHANGE HISTORY**

50694 First released in Issue 4.

50695 **Issue 5**

50696 The following change has been made in this issue for alignment with  
50697 ISO/IEC 9899:1990/Amendment 1:1995 (E):

- 50698 • The SYNOPSIS has been changed to indicate that this function and associated data types are  
50699 now made visible by inclusion of the `<wctype.h>` header rather than `<wchar.h>`.

50700 **NAME**

50701 `wcwidth` — number of column positions of a wide-character code

50702 **SYNOPSIS**

```
50703 xsi #include <wchar.h>
```

```
50704 int wcwidth(wchar_t wc);
```

50705

50706 **DESCRIPTION**

50707 The `wcwidth()` function shall determine the number of column positions required for the wide  
50708 character `wc`. The application shall ensure that the value of `wc` is a character representable as a  
50709 **wchar\_t**, and is a wide-character code corresponding to a valid character in the current locale. |

50710 **RETURN VALUE**

50711 The `wcwidth()` function shall either return 0 (if `wc` is a null wide-character code), or return the  
50712 number of column positions to be occupied by the wide-character code `wc`, or return -1 (if `wc`  
50713 does not correspond to a printable wide-character code). |

50714 **ERRORS**

50715 No errors are defined.

50716 **EXAMPLES**

50717 None.

50718 **APPLICATION USAGE**

50719 This function was removed from the final ISO/IEC 9899:1990/Amendment 1:1995 (E), and the  
50720 return value for a non-printable wide character is not specified.

50721 **RATIONALE**

50722 None.

50723 **FUTURE DIRECTIONS**

50724 None.

50725 **SEE ALSO**

50726 `wcswidth()`, the Base Definitions volume of IEEE Std 1003.1-200x, `<wchar.h>`

50727 **CHANGE HISTORY**

50728 First released as a World-wide Portability Interface in Issue 4. Derived from MSE working draft.

50729 **Issue 6**

50730 The Open Group Corrigendum U021/12 is applied. This function is marked as an extension.

50731 The DESCRIPTION is updated to avoid use of the term “must” for application requirements.



50732 **NAME**

50733 wmemchr — find a wide character in memory

50734 **SYNOPSIS**

50735 #include &lt;wchar.h&gt;

50736 wchar\_t \*wmemchr(const wchar\_t \*ws, wchar\_t wc, size\_t n);

50737 **DESCRIPTION**

50738 cx The functionality described on this reference page is aligned with the ISO C standard. Any  
50739 conflict between the requirements described here and the ISO C standard is unintentional. This  
50740 volume of IEEE Std 1003.1-200x defers to the ISO C standard.

50741 The *wmemchr()* function shall locate the first occurrence of *wc* in the initial *n* wide characters of  
50742 the object pointed to by *ws*. This function shall not be affected by locale and all **wchar\_t** values  
50743 shall be treated identically. The null wide character and **wchar\_t** values not corresponding to  
50744 valid characters shall not be treated specially.

50745 If *n* is zero, the application shall ensure that *ws* is a valid pointer and the function behaves as if  
50746 no valid occurrence of *wc* is found.

50747 **RETURN VALUE**

50748 The *wmemchr()* function shall return a pointer to the located wide character, or a null pointer if  
50749 the wide character does not occur in the object.

50750 **ERRORS**

50751 No errors are defined.

50752 **EXAMPLES**

50753 None.

50754 **APPLICATION USAGE**

50755 None.

50756 **RATIONALE**

50757 None.

50758 **FUTURE DIRECTIONS**

50759 None.

50760 **SEE ALSO**

50761 *wmemcmp()*, *wmemcpy()*, *wmemmove()*, *wmemset()*, the Base Definitions volume of  
50762 IEEE Std 1003.1-200x, <wchar.h>

50763 **CHANGE HISTORY**

50764 First released in Issue 5. Included for alignment with ISO/IEC 9899:1990/Amendment 1:1995  
50765 (E).

50766 **Issue 6**

50767 The DESCRIPTION is updated to avoid use of the term “must” for application requirements.

50768 **NAME**

50769       wmemcmp — compare wide characters in memory

50770 **SYNOPSIS**

50771       #include &lt;wchar.h&gt;

50772       int wmemcmp(const wchar\_t \*ws1, const wchar\_t \*ws2, size\_t n);

50773 **DESCRIPTION**50774 **CX**       The functionality described on this reference page is aligned with the ISO C standard. Any  
50775 conflict between the requirements described here and the ISO C standard is unintentional. This  
50776 volume of IEEE Std 1003.1-200x defers to the ISO C standard.50777       The *wmemcmp()* function shall compare the first *n* wide characters of the object pointed to by  
50778 *ws1* to the first *n* wide characters of the object pointed to by *ws2*. This function shall not be  
50779 affected by locale and all **wchar\_t** values shall be treated identically. The null wide character and  
50780 **wchar\_t** values not corresponding to valid characters shall not be treated specially.50781       If *n* is zero, the application shall ensure that *ws1* and *ws2* are valid pointers, and the function  
50782 shall behave as if the two objects compare equal.50783 **RETURN VALUE**50784       The *wmemcmp()* function shall return an integer greater than, equal to, or less than zero,  
50785 respectively, as the object pointed to by *ws1* is greater than, equal to, or less than the object  
50786 pointed to by *ws2*.50787 **ERRORS**

50788       No errors are defined.

50789 **EXAMPLES**

50790       None.

50791 **APPLICATION USAGE**

50792       None.

50793 **RATIONALE**

50794       None.

50795 **FUTURE DIRECTIONS**

50796       None.

50797 **SEE ALSO**50798       *wmemchr()*, *wmemcpy()*, *wmemmove()*, *wmemset()*, the Base Definitions volume of  
50799 IEEE Std 1003.1-200x, <**wchar.h**>50800 **CHANGE HISTORY**50801       First released in Issue 5. Included for alignment with ISO/IEC 9899:1990/Amendment 1:1995  
50802 (E).50803 **Issue 6**

50804       The DESCRIPTION is updated to avoid use of the term “must” for application requirements.

50805 **NAME**

50806 wmemcpy — copy wide characters in memory

50807 **SYNOPSIS**

50808 #include &lt;wchar.h&gt;

50809 wchar\_t \*wmemcpy(wchar\_t \*restrict ws1, const wchar\_t \*restrict ws2,  
50810 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 *wmemcpy()* function shall copy *n* wide characters from the object pointed to by *ws2* to the  
50816 object pointed to by *ws1*. This function shall not be affected by locale and all **wchar\_t** values  
50817 shall be treated identically. The null wide character and **wchar\_t** values not corresponding to  
50818 valid characters shall not be treated specially. |50819 If *n* is zero, the application shall ensure that *ws1* and *ws2* are valid pointers, and the function  
50820 shall copy zero wide characters. |50821 **RETURN VALUE**50822 The *wmemcpy()* function shall return the value of *ws1*.50823 **ERRORS**

50824 No errors are defined.

50825 **EXAMPLES**

50826 None.

50827 **APPLICATION USAGE**

50828 None.

50829 **RATIONALE**

50830 None.

50831 **FUTURE DIRECTIONS**

50832 None.

50833 **SEE ALSO**50834 *wmemchr()*, *wmemcmp()*, *wmemmove()*, *wmemset()*, the Base Definitions volume of  
50835 IEEE Std 1003.1-200x, <**wchar.h**>50836 **CHANGE HISTORY**50837 First released in Issue 5. Included for alignment with ISO/IEC 9899:1990/Amendment 1:1995  
50838 (E).50839 **Issue 6**

50840 The DESCRIPTION is updated to avoid use of the term “must” for application requirements.

50841 The *wmemcpy()* prototype is updated for alignment with the ISO/IEC 9899:1999 standard.

50842 **NAME**

50843 wmemmove — copy wide characters in memory with overlapping areas

50844 **SYNOPSIS**

50845 #include &lt;wchar.h&gt;

50846 wchar\_t \*wmemmove(wchar\_t \*ws1, const wchar\_t \*ws2, size\_t n);

50847 **DESCRIPTION**

50848 cx The functionality described on this reference page is aligned with the ISO C standard. Any  
50849 conflict between the requirements described here and the ISO C standard is unintentional. This  
50850 volume of IEEE Std 1003.1-200x defers to the ISO C standard.

50851 The *wmemmove()* function shall copy *n* wide characters from the object pointed to by *ws2* to the  
50852 object pointed to by *ws1*. Copying shall take place as if the *n* wide characters from the object  
50853 pointed to by *ws2* are first copied into a temporary array of *n* wide characters that does not  
50854 overlap the objects pointed to by *ws1* or *ws2*, and then the *n* wide characters from the temporary  
50855 array are copied into the object pointed to by *ws1*.

50856 This function shall not be affected by locale and all **wchar\_t** values shall be treated identically.  
50857 The null wide character and **wchar\_t** values not corresponding to valid characters shall not be  
50858 treated specially.

50859 If *n* is zero, the application shall ensure that *ws1* and *ws2* are valid pointers, and the function  
50860 shall copy zero wide characters.

50861 **RETURN VALUE**50862 The *wmemmove()* function shall return the value of *ws1*.50863 **ERRORS**

50864 No errors are defined

50865 **EXAMPLES**

50866 None.

50867 **APPLICATION USAGE**

50868 None.

50869 **RATIONALE**

50870 None.

50871 **FUTURE DIRECTIONS**

50872 None.

50873 **SEE ALSO**

50874 *wmemchr()*, *wmemcmp()*, *wmemcpy()*, *wmemset()*, the Base Definitions volume of  
50875 IEEE Std 1003.1-200x, <**wchar.h**>

50876 **CHANGE HISTORY**

50877 First released in Issue 5. Included for alignment with ISO/IEC 9899:1990/Amendment 1:1995  
50878 (E).

50879 **Issue 6**

50880 The DESCRIPTION is updated to avoid use of the term “must” for application requirements.

50881 **NAME**

50882 wmemset — set wide characters in memory

50883 **SYNOPSIS**

50884 #include &lt;wchar.h&gt;

50885 wchar\_t \*wmemset(wchar\_t \*ws, wchar\_t wc, size\_t n);

50886 **DESCRIPTION**

50887 cx The functionality described on this reference page is aligned with the ISO C standard. Any  
50888 conflict between the requirements described here and the ISO C standard is unintentional. This  
50889 volume of IEEE Std 1003.1-200x defers to the ISO C standard.

50890 The *wmemset()* function shall copy the value of *wc* into each of the first *n* wide characters of the  
50891 object pointed to by *ws*. This function shall not be affected by locale and all **wchar\_t** values shall  
50892 be treated identically. The null wide character and **wchar\_t** values not corresponding to valid  
50893 characters shall not be treated specially.

50894 If *n* is zero, the application shall ensure that *ws* is a valid pointer, and the function shall copy  
50895 zero wide characters.

50896 **RETURN VALUE**50897 The *wmemset()* functions shall return the value of *ws*.50898 **ERRORS**

50899 No errors are defined.

50900 **EXAMPLES**

50901 None.

50902 **APPLICATION USAGE**

50903 None.

50904 **RATIONALE**

50905 None.

50906 **FUTURE DIRECTIONS**

50907 None.

50908 **SEE ALSO**

50909 *wmemchr()*, *wmemcmp()*, *wmemcpy()*, *wmemmove()*, the Base Definitions volume of  
50910 IEEE Std 1003.1-200x, <**wchar.h**>

50911 **CHANGE HISTORY**

50912 First released in Issue 5. Included for alignment with ISO/IEC 9899:1990/Amendment 1:1995  
50913 (E).

50914 **Issue 6**

50915 The DESCRIPTION is updated to avoid use of the term “must” for application requirements.

## 50916 NAME

50917 wordexp, wordfree — perform word expansions

## 50918 SYNOPSIS

50919 #include &lt;wordexp.h&gt;

50920 int wordexp(const char \*restrict words, wordexp\_t \*restrict pwordexp,  
50921 int flags);

50922 void wordfree(wordexp\_t \*pwordexp);

## 50923 DESCRIPTION

50924 The *wordexp()* function shall perform word expansions as described in the Shell and Utilities |  
 50925 volume of IEEE Std 1003.1-200x, Section 2.6, Word Expansions, subject to quoting as in the Shell |  
 50926 and Utilities volume of IEEE Std 1003.1-200x, Section 2.2, Quoting, and place the list of expanded |  
 50927 words into the structure pointed to by *pwordexp*. |

50928 The *words* argument is a pointer to a string containing one or more words to be expanded. The |  
 50929 expansions shall be the same as would be performed by the command line interpreter if *words* |  
 50930 were the part of a command line representing the arguments to a utility. Therefore, the |  
 50931 application shall ensure that *words* does not contain an unquoted <newline> or any of the |  
 50932 unquoted shell special characters `'|', '&', ';', '<', '>'` except in the context of command |  
 50933 substitution as specified in the Shell and Utilities volume of IEEE Std 1003.1-200x, Section 2.6.3, |  
 50934 Command Substitution. It also shall not contain unquoted parentheses or braces, except in the |  
 50935 context of command or variable substitution. The application shall ensure that every member of |  
 50936 *words* which it expects to have expanded by *wordexp()* does not contain an unquoted initial |  
 50937 comment character. The application shall also ensure that any words which it intends to be |  
 50938 ignored (because they begin or continue a comment) are deleted from *words*. If the argument |  
 50939 *words* contains an unquoted comment character (number sign) that is the beginning of a token, |  
 50940 *wordexp()* shall either treat the comment character as a regular character, or interpret it as a |  
 50941 comment indicator and ignore the remainder of *words*.

50942 The structure type **wordexp\_t** is defined in the **<wordexp.h>** header and includes at least the |  
 50943 following members: |

50944

50945

Member Type	Member Name	Description
size_t	<i>we_wordc</i>	Count of words matched by <i>words</i> .
char **	<i>we_wordv</i>	Pointer to list of expanded words.
size_t	<i>we_offs</i>	Slots to reserve at the beginning of <i>pwordexp-&gt;we_wordv</i> .

50946

50947

50948

50949 The *wordexp()* function shall store the number of generated words into *pwordexp->we\_wordc* and |  
 50950 a pointer to a list of pointers to words in *pwordexp->we\_wordv*. Each individual field created |  
 50951 during field splitting (see the Shell and Utilities volume of IEEE Std 1003.1-200x, Section 2.6.5, |  
 50952 Field Splitting) or pathname expansion (see the Shell and Utilities volume of |  
 50953 IEEE Std 1003.1-200x, Section 2.6.6, Pathname Expansion) shall be a separate word in the |  
 50954 *pwordexp->we\_wordv* list. The words shall be in order as described in the Shell and Utilities |  
 50955 volume of IEEE Std 1003.1-200x, Section 2.6, Word Expansions. The first pointer after the last |  
 50956 word pointer shall be a null pointer. The expansion of special parameters described in the Shell |  
 50957 and Utilities volume of IEEE Std 1003.1-200x, Section 2.5.2, Special Parameters is unspecified. |

50958 It is the caller's responsibility to allocate the storage pointed to by *pwordexp*. The *wordexp()* |  
 50959 function shall allocate other space as needed, including memory pointed to by *pwordexp-* |  
 50960 *>we\_wordv*. The *wordfree()* function frees any memory associated with *pwordexp* from a previous |  
 50961 call to *wordexp()*.

50962 The *flags* argument is used to control the behavior of *wordexp()*. The value of *flags* is the  
50963 bitwise-inclusive OR of zero or more of the following constants, which are defined in  
50964 **<wordexp.h>**:

50965 **WRDE\_APPEND** Append words generated to the ones from a previous call to *wordexp()*.

50966 **WRDE\_DOOFFS** Make use of *pwordexp->we\_offs*. If this flag is set, *pwordexp->we\_offs* is used  
50967 to specify how many null pointers to add to the beginning of *pwordexp->we\_wordv*. In other words, *pwordexp->we\_wordv*  
50968 shall point to *pwordexp->we\_offs* null pointers, followed by *pwordexp->we\_wordc* word pointers,  
50970 followed by a null pointer.

50971 **WRDE\_NOCMD** If the implementation supports the utilities defined in the Shell and  
50972 Utilities volume of IEEE Std 1003.1-200x, fail if command substitution, as  
50973 specified in the Shell and Utilities volume of IEEE Std 1003.1-200x,  
50974 Section 2.6.3, Command Substitution, is requested.

50975 **WRDE\_REUSE** The *pwordexp* argument was passed to a previous successful call to  
50976 *wordexp()*, and has not been passed to *wordfree()*. The result shall be the  
50977 same as if the application had called *wordfree()* and then called *wordexp()*  
50978 without **WRDE\_REUSE**.

50979 **WRDE\_SHOWERR** Do not redirect *stderr* to **/dev/null**.

50980 **WRDE\_UNDEF** Report error on an attempt to expand an undefined shell variable.

50981 The **WRDE\_APPEND** flag can be used to append a new set of words to those generated by a  
50982 previous call to *wordexp()*. The following rules apply to applications when two or more calls to  
50983 *wordexp()* are made with the same value of *pwordexp* and without intervening calls to *wordfree()*:

- 50984 1. The first such call shall not set **WRDE\_APPEND**. All subsequent calls shall set it.
- 50985 2. All of the calls shall set **WRDE\_DOOFFS**, or all shall not set it.
- 50986 3. After the second and each subsequent call, *pwordexp->we\_wordv* shall point to a list  
50987 containing the following:
  - 50988 a. Zero or more null pointers, as specified by **WRDE\_DOOFFS** and *pwordexp->we\_offs*
  - 50989 b. Pointers to the words that were in the *pwordexp->we\_wordv* list before the call, in the  
50990 same order as before
  - 50991 c. Pointers to the new words generated by the latest call, in the specified order
- 50992 4. The count returned in *pwordexp->we\_wordc* shall be the total number of words from all of  
50993 the calls.
- 50994 5. The application can change any of the fields after a call to *wordexp()*, but if it does it shall  
50995 reset them to the original value before a subsequent call, using the same *pwordexp* value, to  
50996 *wordfree()* or *wordexp()* with the **WRDE\_APPEND** or **WRDE\_REUSE** flag.

50997 If the implementation supports the utilities defined in the Shell and Utilities volume of  
50998 IEEE Std 1003.1-200x, and *words* contains an unquoted character—<newline>, '|', '&', ';',  
50999 '<', '>', '(', ')', '{', '}'—in an inappropriate context, *wordexp()* shall fail, and the number  
51000 of expanded words shall be 0.

51001 Unless **WRDE\_SHOWERR** is set in *flags*, *wordexp()* shall redirect *stderr* to **/dev/null** for any  
51002 utilities executed as a result of command substitution while expanding *words*. If  
51003 **WRDE\_SHOWERR** is set, *wordexp()* may write messages to *stderr* if syntax errors are detected  
51004 while expanding *words*.

51005 The application shall ensure that if WRDE\_DOOFFS is set, then *pwordexp->we\_offs* has the same  
51006 value for each *wordexp()* call and *wordfree()* call using a given *pwordexp*.

51007 The following constants are defined as error return values:

51008 WRDE\_BADCHAR One of the unquoted characters—<newline>, '|', '&', ';', '<', '>',  
51009 '(', ')', '{', '}'—appears in *words* in an inappropriate context.

51010 WRDE\_BADVAL Reference to undefined shell variable when WRDE\_UNDEF is set in *flags*.

51011 WRDE\_CMDSUB Command substitution requested when WRDE\_NOCMD was set in *flags*.

51012 WRDE\_NOSPACE Attempt to allocate memory failed.

51013 WRDE\_SYNTAX Shell syntax error, such as unbalanced parentheses or unterminated  
51014 string.

#### 51015 RETURN VALUE

51016 Upon successful completion, *wordexp()* shall return 0. Otherwise, a non-zero value, as described  
51017 in <**wordexp.h**>, shall be returned to indicate an error. If *wordexp()* returns the value  
51018 WRDE\_NOSPACE, then *pwordexp->we\_wordc* and *pwordexp->we\_wordv* shall be updated to  
51019 reflect any words that were successfully expanded. In other cases, they shall not be modified.

51020 The *wordfree()* function shall not return a value.

#### 51021 ERRORS

51022 No errors are defined.

#### 51023 EXAMPLES

51024 None.

#### 51025 APPLICATION USAGE

51026 The *wordexp()* function is intended to be used by an application that wants to do all of the shell's  
51027 expansions on a word or words obtained from a user. For example, if the application prompts  
51028 for a filename (or list of filenames) and then uses *wordexp()* to process the input, the user could  
51029 respond with anything that would be valid as input to the shell.

51030 The WRDE\_NOCMD flag is provided for applications that, for security or other reasons, want to  
51031 prevent a user from executing shell commands. Disallowing unquoted shell special characters  
51032 also prevents unwanted side effects, such as executing a command or writing a file.

#### 51033 RATIONALE

51034 This function was included as an alternative to *glob()*. There had been continuing controversy  
51035 over exactly what features should be included in *glob()*. It is hoped that by providing *wordexp()*  
51036 (which provides all of the shell word expansions, but which may be slow to execute) and *glob()* |  
51037 (which is faster, but which only performs pathname expansion, without tilde or parameter |  
51038 expansion) this will satisfy the majority of applications.

51039 While *wordexp()* could be implemented entirely as a library routine, it is expected that most  
51040 implementations run a shell in a subprocess to do the expansion.

51041 Two different approaches have been proposed for how the required information might be  
51042 presented to the shell and the results returned. They are presented here as examples.

51043 One proposal is to extend the *echo* utility by adding a **-q** option. This option would cause *echo*  
51044 to add a backslash before each backslash and <blank> that occurs within an argument. The  
51045 *wordexp()* function could then invoke the shell as follows:

```
51046 (void) strcpy(buffer, "echo -q");
51047 (void) strcat(buffer, words);
51048 if ((flags & WRDE_SHOWERR) == 0)
```



```
51049         (void) strcat(buffer, "2>/dev/null");
51050         f = popen(buffer, "r");
```

51051 The *wordexp()* function would read the resulting output, remove unquoted backslashes, and  
 51052 break into words at unquoted <blank>s. If the WRDE\_NOCMD flag was set, *wordexp()* would  
 51053 have to scan *words* before starting the subshell to make sure that there would be no command  
 51054 substitution. In any case, it would have to scan *words* for unquoted special characters.

51055 Another proposal is to add the following options to *sh*:

51056 **-w wordlist**

51057 This option provides a wordlist expansion service to applications. The words in *wordlist*  
 51058 shall be expanded and the following written to standard output:

- 51059 1. The count of the number of words after expansion, in decimal, followed by a null byte
- 51060 2. The number of bytes needed to represent the expanded words (not including null  
 51061 separators), in decimal, followed by a null byte
- 51062 3. The expanded words, each terminated by a null byte

51063 If an error is encountered during word expansion, *sh* exits with a non-zero status after  
 51064 writing the former to report any words successfully expanded

51065 **-P** Run in “protected” mode. If specified with the **-w** option, no command substitution shall  
 51066 be performed.

51067 With these options, *wordexp()* could be implemented fairly simply by creating a subprocess  
 51068 using *fork()* and executing *sh* using the line:

```
51069 execl(<shell path>, "sh", "-P", "-w", words, (char *)0);
```

51070 after directing standard error to **/dev/null**.

51071 It seemed objectionable for a library routine to write messages to standard error, unless  
 51072 explicitly requested, so *wordexp()* is required to redirect standard error to **/dev/null** to ensure  
 51073 that no messages are generated, even for commands executed for command substitution. The  
 51074 WRDE\_SHOWERR flag can be specified to request that error messages be written.

51075 The WRDE\_REUSE flag allows the implementation to avoid the expense of freeing and  
 51076 reallocating memory, if that is possible. A minimal implementation can call *wordfree()* when  
 51077 WRDE\_REUSE is set.

#### 51078 FUTURE DIRECTIONS

51079 None.

#### 51080 SEE ALSO

51081 *fnmatch()*, *glob()*, the Base Definitions volume of IEEE Std 1003.1-200x, **<wordexp.h>**, the Shell  
 51082 and Utilities volume of IEEE Std 1003.1-200x

#### 51083 CHANGE HISTORY

51084 First released in Issue 4. Derived from the ISO POSIX-2 standard.

#### 51085 Issue 5

51086 Moved from POSIX2 C-language Binding to BASE.

#### 51087 Issue 6

51088 The DESCRIPTION is updated to avoid use of the term “must” for application requirements.

51089 The **restrict** keyword is added to the *wordexp()* prototype for alignment with the  
 51090 ISO/IEC 9899:1999 standard.

51091 **NAME**

51092        `wprintf` — print formatted wide-character output

51093 **SYNOPSIS**

51094        `#include <stdio.h>`

51095        `#include <wchar.h>`

51096        `int wprintf(const wchar_t *restrict format, ...);` |

51097 **DESCRIPTION** |

51098        Refer to *fwprintf()*.

## 51099 NAME

51100 pwrite, write — write on a file

## 51101 SYNOPSIS

51102 #include &lt;unistd.h&gt;

51103 XSI ssize\_t pwrite(int *fildev*, const void \**buf*, size\_t *nbyte*,  
51104 off\_t *offset*);51105 ssize\_t write(int *fildev*, const void \**buf*, size\_t *nbyte*);

## 51106 DESCRIPTION

51107 The *write()* function shall attempt to write *nbyte* bytes from the buffer pointed to by *buf* to the  
51108 file associated with the open file descriptor, *fildev*.51109 Before any action described below is taken, and if *nbyte* is zero and the file is a regular file, the  
51110 *write()* function may detect and return errors as described below. In the absence of errors, or if  
51111 error detection is not performed, the *write()* function shall return zero and have no other results.  
51112 If *nbyte* is zero and the file is not a regular file, the results are unspecified.51113 On a regular file or other file capable of seeking, the actual writing of data shall proceed from the  
51114 position in the file indicated by the file offset associated with *fildev*. Before successful return  
51115 from *write()*, the file offset shall be incremented by the number of bytes actually written. On a  
51116 regular file, if this incremented file offset is greater than the length of the file, the length of the  
51117 file shall be set to this file offset.51118 On a file not capable of seeking, writing shall always take place starting at the current position.  
51119 The value of a file offset associated with such a device is undefined.51120 If the O\_APPEND flag of the file status flags is set, the file offset shall be set to the end of the file  
51121 prior to each write and no intervening file modification operation shall occur between changing  
51122 the file offset and the write operation.51123 XSI If a *write()* requests that more bytes be written than there is room for (for example, the process'  
51124 file size limit or the physical end of a medium), only as many bytes as there is room for shall be  
51125 written. For example, suppose there is space for 20 bytes more in a file before reaching a limit. A  
51126 write of 512 bytes will return 20. The next write of a non-zero number of bytes would give a  
51127 failure return (except as noted below).51128 XSI If the request would cause the file size to exceed the soft file size limit for the process and there  
51129 is no room for any bytes to be written, the request shall fail and the implementation shall  
51130 generate the SIGXFSZ signal for the thread.51131 If *write()* is interrupted by a signal before it writes any data, it shall return  $-1$  with *errno* set to  
51132 [EINTR].51133 If *write()* is interrupted by a signal after it successfully writes some data, it shall return the  
51134 number of bytes written.51135 If the value of *nbyte* is greater than {SSIZE\_MAX}, the result is implementation-defined.51136 After a *write()* to a regular file has successfully returned:

- 51137
- Any successful *read()* from each byte position in the file that was modified by that write shall  
51138 return the data specified by the *write()* for that position until such byte positions are again  
51139 modified.
  - Any subsequent successful *write()* to the same byte position in the file shall overwrite that  
51140 file data.  
51141

51142 Write requests to a pipe or FIFO shall be handled in the same way as a regular file with the  
51143 following exceptions:

51144 • There is no file offset associated with a pipe, hence each write request shall append to the  
51145 end of the pipe.

51146 • Write requests of {PIPE\_BUF} bytes or less shall not be interleaved with data from other  
51147 processes doing writes on the same pipe. Writes of greater than {PIPE\_BUF} bytes may have  
51148 data interleaved, on arbitrary boundaries, with writes by other processes, whether or not the  
51149 O\_NONBLOCK flag of the file status flags is set.

51150 • If the O\_NONBLOCK flag is clear, a write request may cause the thread to block, but on  
51151 normal completion it shall return *nbyte*.

51152 • If the O\_NONBLOCK flag is set, *write()* requests shall be handled differently, in the  
51153 following ways:

51154 — The *write()* function shall not block the thread.

51155 — A write request for {PIPE\_BUF} or fewer bytes shall have the following effect: if there is  
51156 sufficient space available in the pipe, *write()* shall transfer all the data and return the  
51157 number of bytes requested. Otherwise, *write()* shall transfer no data and return  $-1$  with  
51158 *errno* set to [EAGAIN].

51159 — A write request for more than {PIPE\_BUF} bytes shall cause one of the following:

51160 — When at least one byte can be written, transfer what it can and return the number of  
51161 bytes written. When all data previously written to the pipe is read, it shall transfer at  
51162 least {PIPE\_BUF} bytes.

51163 — When no data can be written, transfer no data, and return  $-1$  with *errno* set to  
51164 [EAGAIN].

51165 When attempting to write to a file descriptor (other than a pipe or FIFO) that supports non-  
51166 blocking writes and cannot accept the data immediately:

51167 • If the O\_NONBLOCK flag is clear, *write()* shall block the calling thread until the data can be  
51168 accepted.

51169 • If the O\_NONBLOCK flag is set, *write()* shall not block the thread. If some data can be  
51170 written without blocking the thread, *write()* shall write what it can and return the number of  
51171 bytes written. Otherwise, it shall return  $-1$  and set *errno* to [EAGAIN].

51172 Upon successful completion, where *nbyte* is greater than 0, *write()* shall mark for update the  
51173 *st\_ctime* and *st\_mtime* fields of the file, and if the file is a regular file, the S\_ISUID and S\_ISGID  
51174 bits of the file mode may be cleared.

51175 For regular files, no data transfer shall occur past the offset maximum established in the open  
51176 file description associated with *fildes*.

51177 If *fildes* refers to a socket, *write()* shall be equivalent to *send()* with no flags set.

51178 SIO If the O\_DSYNC bit has been set, write I/O operations on the file descriptor shall complete as  
51179 defined by synchronized I/O data integrity completion.

51180 If the O\_SYNC bit has been set, write I/O operations on the file descriptor shall complete as  
51181 defined by synchronized I/O file integrity completion.

51182 SHM If *fildes* refers to a shared memory object, the result of the *write()* function is unspecified.

51183 TYM If *fildes* refers to a typed memory object, the result of the *write()* function is unspecified.

51184 XSR If *fildev* refers to a STREAM, the operation of *write()* shall be determined by the values of the  
 51185 minimum and maximum *nbyte* range (packet size) accepted by the STREAM. These values are  
 51186 determined by the topmost STREAM module. If *nbyte* falls within the packet size range, *nbyte*  
 51187 bytes shall be written. If *nbyte* does not fall within the range and the minimum packet size value  
 51188 is 0, *write()* shall break the buffer into maximum packet size segments prior to sending the data  
 51189 downstream (the last segment may contain less than the maximum packet size). If *nbyte* does not  
 51190 fall within the range and the minimum value is non-zero, *write()* shall fail with *errno* set to  
 51191 [ERANGE]. Writing a zero-length buffer (*nbyte* is 0) to a STREAMS device sends 0 bytes with 0  
 51192 returned. However, writing a zero-length buffer to a STREAMS-based pipe or FIFO sends no  
 51193 message and 0 is returned. The process may issue *I\_SWROPT ioctl()* to enable zero-length  
 51194 messages to be sent across the pipe or FIFO.

51195 When writing to a STREAM, data messages are created with a priority band of 0. When writing  
 51196 to a STREAM that is not a pipe or FIFO:

- 51197 • If O\_NONBLOCK is clear, and the STREAM cannot accept data (the STREAM write queue is  
 51198 full due to internal flow control conditions), *write()* shall block until data can be accepted.

- 51199 • If O\_NONBLOCK is set and the STREAM cannot accept data, *write()* shall return -1 and set  
 51200 *errno* to [EAGAIN].

- 51201 • If O\_NONBLOCK is set and part of the buffer has been written while a condition in which  
 51202 the STREAM cannot accept additional data occurs, *write()* shall terminate and return the  
 51203 number of bytes written.

51204 In addition, *write()* shall fail if the STREAM head has processed an asynchronous error before  
 51205 the call. In this case, the value of *errno* does not reflect the result of *write()*, but reflects the prior  
 51206 error.

51207 XSI The *pwrite()* function shall be equivalent to *write()*, except that it writes into a given position  
 51208 without changing the file pointer. The first three arguments to *pwrite()* are the same as *write()*  
 51209 with the addition of a fourth argument offset for the desired position inside the file.

#### 51210 RETURN VALUE

51211 XSI Upon successful completion, *write()* and *pwrite()* shall return the number of bytes actually  
 51212 written to the file associated with *fildev*. This number shall never be greater than *nbyte*.  
 51213 Otherwise, -1 shall be returned and *errno* set to indicate the error.

#### 51214 ERRORS

51215 XSI The *write()* and *pwrite()* functions shall fail if:

51216 [EAGAIN] The O\_NONBLOCK flag is set for the file descriptor and the thread would be  
 51217 delayed in the *write()* operation.

51218 [EBADF] The *fildev* argument is not a valid file descriptor open for writing.

51219 [EFBIG] An attempt was made to write a file that exceeds the implementation-defined  
 51220 XSI maximum file size or the process' file size limit, and there was no room for  
 51221 any bytes to be written.

51222 [EFBIG] The file is a regular file, *nbyte* is greater than 0, and the starting position is  
 51223 greater than or equal to the offset maximum established in the open file  
 51224 description associated with *fildev*.

51225 [EINTR] The write operation was terminated due to the receipt of a signal, and no data  
 51226 was transferred.

51227 [EIO] The process is a member of a background process group attempting to write  
 51228 to its controlling terminal, TOSTOP is set, the process is neither ignoring nor

51229		blocking SIGTTOU, and the process group of the process is orphaned. This error may also be returned under implementation-defined conditions.
51230		
51231	[ENOSPC]	There was no free space remaining on the device containing the file.
51232	[EPIPE]	An attempt is made to write to a pipe or FIFO that is not open for reading by any process, or that only has one end open. A SIGPIPE signal shall also be sent to the thread.
51233		
51234		
51235 XSR	[ERANGE]	The transfer request size was outside the range supported by the STREAMS file associated with <i>fildev</i> .
51236		
51237		The <i>write()</i> function shall fail if:
51238	[EAGAIN] or [EWOULDBLOCK]	
51239		The file descriptor is for a socket, is marked O_NONBLOCK, and write would block.
51240		
51241	[ECONNRESET]	A write was attempted on a socket that is not connected.
51242	[EPIPE]	A write was attempted on a socket that is shut down for writing, or is no longer connected. In the latter case, if the socket is of type SOCK_STREAM, the SIGPIPE signal is generated to the calling process.
51243		
51244		
51245 XSI		The <i>write()</i> and <i>pwrite()</i> functions may fail if:
51246 XSR	[EINVAL]	The STREAM or multiplexer referenced by <i>fildev</i> is linked (directly or indirectly) downstream from a multiplexer.
51247		
51248	[EIO]	A physical I/O error has occurred.
51249	[ENOBUFS]	Insufficient resources were available in the system to perform the operation.
51250	[ENXIO]	A request was made of a nonexistent device, or the request was outside the capabilities of the device.
51251		
51252 XSR	[ENXIO]	A hangup occurred on the STREAM being written to.
51253 XSR		A write to a STREAMS file may fail if an error message has been received at the STREAM head. In this case, <i>errno</i> is set to the value included in the error message.
51254		
51255		The <i>write()</i> function may fail if:
51256	[EACCES]	A write was attempted on a socket and the calling process does not have appropriate privileges.
51257		
51258	[ENETDOWN]	A write was attempted on a socket and the local network interface used to reach the destination is down.
51259		
51260	[ENETUNREACH]	A write was attempted on a socket and no route to the network is present.
51261 XSI		The <i>pwrite()</i> function shall fail and the file pointer remain unchanged if:
51262 XSI	[EINVAL]	The <i>offset</i> argument is invalid. The value is negative.
51263 XSI	[ESPIPE]	<i>fildev</i> is associated with a pipe or FIFO.

51264 **EXAMPLES**51265 **Writing from a Buffer**

51266 The following example writes data from the buffer pointed to by *buf* to the file associated with  
51267 the file descriptor *fd*.

```
51268 #include <sys/types.h>
51269 #include <string.h>
51270 ...
51271 char buf[20];
51272 size_t nbytes;
51273 ssize_t bytes_written;
51274 int fd;
51275 ...
51276 strcpy(buf, "This is a test\n");
51277 nbytes = strlen(buf);

51278 bytes_written = write(fd, buf, nbytes);
51279 ...
```

51280 **APPLICATION USAGE**

51281 None.

51282 **RATIONALE**

51283 See also the RATIONALE section in *read()*.

51284 An attempt to write to a pipe or FIFO has several major characteristics:

- 51285 • *Atomic/non-atomic*: A write is atomic if the whole amount written in one operation is not  
51286 interleaved with data from any other process. This is useful when there are multiple writers  
51287 sending data to a single reader. Applications need to know how large a write request can be  
51288 expected to be performed atomically. This maximum is called {PIPE\_BUF}. This volume of  
51289 IEEE Std 1003.1-200x does not say whether write requests for more than {PIPE\_BUF} bytes  
51290 are atomic, but requires that writes of {PIPE\_BUF} or fewer bytes shall be atomic.
- 51291 • *Blocking/immediate*: Blocking is only possible with O\_NONBLOCK clear. If there is enough  
51292 space for all the data requested to be written immediately, the implementation should do so.  
51293 Otherwise, the process may block; that is, pause until enough space is available for writing.  
51294 The effective size of a pipe or FIFO (the maximum amount that can be written in one  
51295 operation without blocking) may vary dynamically, depending on the implementation, so it  
51296 is not possible to specify a fixed value for it.

- 51297 • *Complete/partial/deferred*: A write request:

```
51298 int fildes;
51299 size_t nbyte;
51300 ssize_t ret;
51301 char *buf;

51302 ret = write(fildes, buf, nbyte);
```

51303 may return:

51304 complete     *ret*=*nbyte*

51305 partial       *ret*<*nbyte*

51306 This shall never happen if *nbyte*≤{PIPE\_BUF}. If it does happen (with  
51307 *nbyte*>{PIPE\_BUF}), this volume of IEEE Std 1003.1-200x does not guarantee

51308 atomicity, even if  $ret \leq \{PIPE\_BUF\}$ , because atomicity is guaranteed according  
51309 to the amount *requested*, not the amount *written*.

51310 deferred:  $ret = -1$ ,  $errno = [EAGAIN]$

51311 This error indicates that a later request may succeed. It does not indicate that it  
51312 *shall* succeed, even if  $nbyte \leq \{PIPE\_BUF\}$ , because if no process reads from the  
51313 pipe or FIFO, the write never succeeds. An application could usefully count the  
51314 number of times  $[EAGAIN]$  is caused by a particular value of  
51315  $nbyte > \{PIPE\_BUF\}$  and perhaps do later writes with a smaller value, on the  
51316 assumption that the effective size of the pipe may have decreased.

51317 Partial and deferred writes are only possible with `O_NONBLOCK` set.

51318 The relations of these properties are shown in the following tables:

51319

51320

51321

51322

51323

51324

Write to a Pipe or FIFO with <code>O_NONBLOCK</code> clear			
Immediately Writable:	None	Some	<i>nbyte</i>
$nbyte \leq \{PIPE\_BUF\}$	Atomic blocking <i>nbyte</i>	Atomic blocking <i>nbyte</i>	Atomic immediate <i>nbyte</i>
$nbyte > \{PIPE\_BUF\}$	Blocking <i>nbyte</i>	Blocking <i>nbyte</i>	Blocking <i>nbyte</i>

51325 If the `O_NONBLOCK` flag is clear, a write request shall block if the amount writable  
51326 immediately is less than that requested. If the flag is set (by `fcntl()`), a write request shall never  
51327 block.

51328

51329

51330

51331

51332

51333

Write to a Pipe or FIFO with <code>O_NONBLOCK</code> set			
Immediately Writable:	None	Some	<i>nbyte</i>
$nbyte \leq \{PIPE\_BUF\}$	-1, $[EAGAIN]$	-1, $[EAGAIN]$	Atomic <i>nbyte</i>
$nbyte > \{PIPE\_BUF\}$	-1, $[EAGAIN]$	$<nbyte$ or -1, $[EAGAIN]$	$\leq nbyte$ or -1, $[EAGAIN]$

51334 There is no exception regarding partial writes when `O_NONBLOCK` is set. With the exception  
51335 of writing to an empty pipe, this volume of IEEE Std 1003.1-200x does not specify exactly when a  
51336 partial write is performed since that would require specifying internal details of the  
51337 implementation. Every application should be prepared to handle partial writes when  
51338 `O_NONBLOCK` is set and the requested amount is greater than  $\{PIPE\_BUF\}$ , just as every  
51339 application should be prepared to handle partial writes on other kinds of file descriptors.

51340 The intent of forcing writing at least one byte if any can be written is to assure that each write  
51341 makes progress if there is any room in the pipe. If the pipe is empty,  $\{PIPE\_BUF\}$  bytes must be  
51342 written; if not, at least some progress must have been made.

51343 Where this volume of IEEE Std 1003.1-200x requires -1 to be returned and *errno* set to  
51344  $[EAGAIN]$ , most historical implementations return zero (with the `O_NDELAY` flag set, which is  
51345 the historical predecessor of `O_NONBLOCK`, but is not itself in this volume of  
51346 IEEE Std 1003.1-200x). The error indications in this volume of IEEE Std 1003.1-200x were chosen  
51347 so that an application can distinguish these cases from end-of-file. While `write()` cannot receive  
51348 an indication of end-of-file, `read()` can, and the two functions have similar return values. Also,  
51349 some existing systems (for example, Eighth Edition) permit a write of zero bytes to mean that  
51350 the reader should get an end-of-file indication; for those systems, a return value of zero from  
51351 `write()` indicates a successful write of an end-of-file indication.



51352 Implementations are allowed, but not required, to perform error checking for *write()* requests of  
51353 zero bytes.

51354 The concept of a {PIPE\_MAX} limit (indicating the maximum number of bytes that can be  
51355 written to a pipe in a single operation) was considered, but rejected, because this concept would  
51356 unnecessarily limit application writing.

51357 See also the discussion of O\_NONBLOCK in *read()*.

51358 Writes can be serialized with respect to other reads and writes. If a *read()* of file data can be  
51359 proven (by any means) to occur after a *write()* of the data, it must reflect that *write()*, even if the  
51360 calls are made by different processes. A similar requirement applies to multiple write operations  
51361 to the same file position. This is needed to guarantee the propagation of data from *write()* calls  
51362 to subsequent *read()* calls. This requirement is particularly significant for networked file  
51363 systems, where some caching schemes violate these semantics.

51364 Note that this is specified in terms of *read()* and *write()*. The XSI extensions *readv()* and *writv()* |  
51365 also obey these semantics. A new “high-performance” write analog that did not follow these |  
51366 serialization requirements would also be permitted by this wording. This volume of  
51367 IEEE Std 1003.1-200x is also silent about any effects of application-level caching (such as that  
51368 done by *stdio*).

51369 This volume of IEEE Std 1003.1-200x does not specify the value of the file offset after an error is  
51370 returned; there are too many cases. For programming errors, such as [EBADF], the concept is  
51371 meaningless since no file is involved. For errors that are detected immediately, such as  
51372 [EAGAIN], clearly the pointer should not change. After an interrupt or hardware error, however,  
51373 an updated value would be very useful and is the behavior of many implementations.

51374 This volume of IEEE Std 1003.1-200x does not specify behavior of concurrent writes to a file  
51375 from multiple processes. Applications should use some form of concurrency control.

#### 51376 FUTURE DIRECTIONS

51377 None.

#### 51378 SEE ALSO

51379 *chmod()*, *creat()*, *dup()*, *fcntl()*, *getrlimit()*, *lseek()*, *open()*, *pipe()*, *ulimit()*, *writv()*, the Base |  
51380 Definitions volume of IEEE Std 1003.1-200x, <limits.h>, <stropts.h>, <sys/uio.h>, <unistd.h>

#### 51381 CHANGE HISTORY

51382 First released in Issue 1. Derived from Issue 1 of the SVID.

#### 51383 Issue 5

51384 The DESCRIPTION is updated for alignment with the POSIX Realtime Extension and the POSIX  
51385 Threads Extension.

51386 Large File Summit extensions are added.

51387 The *pwrite()* function is added.

#### 51388 Issue 6

51389 The DESCRIPTION states that the *write()* function does not block the thread. Previously this  
51390 said “process” rather than “thread”.

51391 The DESCRIPTION and ERRORS sections are updated so that references to STREAMS are  
51392 marked as part of the XSI STREAMS Option Group.

51393 The following new requirements on POSIX implementations derive from alignment with the  
51394 Single UNIX Specification:

- 51395 • The DESCRIPTION now states that if *write()* is interrupted by a signal after it has  
51396 successfully written some data, it returns the number of bytes written. In earlier versions of  
51397 this volume of IEEE Std 1003.1-200x, it was optional whether *write()* returned the number of  
51398 bytes written, or whether it returned  $-1$  with *errno* set to [EINTR]. This is a FIPS requirement.
- 51399 • The following changes are made to support large files:
- 51400 — For regular files, no data transfer occurs past the offset maximum established in the open  
51401 file description associated with the *files*.
  - 51402 — A second [EFBIG] error condition is added.
- 51403 • The [EIO] error condition is added.
- 51404 • The [EPIPE] error condition is added for when a pipe has only one end open.
- 51405 • The [ENXIO] optional error condition is added.
- 51406 Text referring to sockets is added to the DESCRIPTION.
- 51407 The following changes were made to align with the IEEE P1003.1a draft standard:
- 51408 • The effect of reading zero bytes is clarified.
- 51409 The DESCRIPTION is updated for alignment with IEEE Std 1003.1j-2000 by specifying that  
51410 *write()* results are unspecified for typed memory objects.
- 51411 The following error conditions are added for operations on sockets: [EAGAIN],  
51412 [EWOULDBLOCK], [ECONNRESET], [ENOTCONN], and [EPIPE].
- 51413 The [EIO] error is changed to “may fail”.
- 51414 The [ENOBUFFS] error is added for sockets.
- 51415 The following error conditions are added for operations on sockets: [EACCES], [ENETDOWN],  
51416 and [ENETUNREACH].
- 51417 The *writenv()* function is split out into a separate reference page.

51418 **NAME**

51419 writev — write a vector

51420 **SYNOPSIS**51421 XSI `#include <sys/uio.h>`51422 `ssize_t writev(int fildev, const struct iovec *iov, int iovcnt);`

51423

51424 **DESCRIPTION**

51425 The `writev()` function shall be equivalent to `write()`, except as described below. The `writev()`  
 51426 function shall gather output data from the `iovcnt` buffers specified by the members of the `iov`  
 51427 array: `iov[0]`, `iov[1]`, ..., `iov[iovcnt-1]`. The `iovcnt` argument is valid if greater than 0 and less than  
 51428 or equal to `{IOV_MAX}`, defined in `<limits.h>`.

51429 Each `iovec` entry specifies the base address and length of an area in memory from which data  
 51430 should be written. The `writev()` function shall always write a complete area before proceeding to  
 51431 the next.

51432 If `fildev` refers to a regular file and all of the `iov_len` members in the array pointed to by `iov` are 0,  
 51433 `writev()` shall return 0 and have no other effect. For other file types, the behavior is unspecified.

51434 If the sum of the `iov_len` values is greater than `{SSIZE_MAX}`, the operation shall fail and no data  
 51435 shall be transferred.

51436 **RETURN VALUE**

51437 Upon successful completion, `writev()` shall return the number of bytes actually written.  
 51438 Otherwise, it shall return a value of `-1`, the file-pointer shall remain unchanged, and `errno` shall  
 51439 be set to indicate an error.

51440 **ERRORS**51441 Refer to `write()`.51442 In addition, the `writev()` function shall fail if:51443 [EINVAL] The sum of the `iov_len` values in the `iov` array would overflow an `ssize_t`.51444 The `writev()` function may fail and set `errno` to:51445 [EINVAL] The `iovcnt` argument was less than or equal to 0, or greater than `{IOV_MAX}`.51446 **EXAMPLES**51447 **Writing Data from an Array**

51448 The following example writes data from the buffers specified by members of the `iov` array to the  
 51449 file associated with the file descriptor `fd`.

51450 `#include <sys/types.h>`51451 `#include <sys/uio.h>`51452 `#include <unistd.h>`51453 `...`51454 `ssize_t bytes_written;`51455 `int fd;`51456 `char *buf0 = "short string\n";`51457 `char *buf1 = "This is a longer string\n";`51458 `char *buf2 = "This is the longest string in this example\n";`51459 `int iovcnt;`51460 `struct iovec iov[3];`

```
51461     iov[0].iov_base = buf0;
51462     iov[0].iov_len = strlen(buf0);
51463     iov[1].iov_base = buf1;
51464     iov[1].iov_len = strlen(buf1);
51465     iov[2].iov_base = buf2;
51466     iov[2].iov_len = strlen(buf2);
51467     ...
51468     iovcnt = sizeof(iov) / sizeof(struct iovec);

51469     bytes_written = writev(fd, iov, iovcnt);
51470     ...
```

**51471 APPLICATION USAGE**

51472 None.

**51473 RATIONALE**

51474 Refer to *write()*.

**51475 FUTURE DIRECTIONS**

51476 None.

**51477 SEE ALSO**

51478 *readv()*, *write()*, the Base Definitions volume of IEEE Std 1003.1-200x, <limits.h>, <sys/uio.h>

**51479 CHANGE HISTORY**

51480 First released in Issue 4, Version 2.

**51481 Issue 6**

51482 Split out from the *write()* reference page.

51483 **NAME**

51484           wscanf — convert formatted wide-character input

51485 **SYNOPSIS**

51486           #include &lt;stdio.h&gt;

51487           #include &lt;wchar.h&gt;

51488           int wscanf(const wchar\_t \*restrict *format*, ... );51489 **DESCRIPTION**51490           Refer to *fwscanf()*.

51491 **NAME**

51492        y0, y1, yn — Bessel functions of the second kind

51493 **SYNOPSIS**

```
51494 xSI      #include <math.h>
51495          double y0(double x);
51496          double y1(double x);
51497          double yn(int n, double x);
51498
```

51499 **DESCRIPTION**

51500        The *y0()*, *y1()*, and *yn()* functions shall compute Bessel functions of *x* of the second kind of  
 51501        orders 0, 1, and *n*, respectively.

51502        An application wishing to check for error situations should set *errno* to zero and call  
 51503        *feclearexcept*(FE\_ALL\_EXCEPT) before calling these functions. On return, if *errno* is non-zero or  
 51504        *fetestexcept*(FE\_INVALID | FE\_DIVBYZERO | FE\_OVERFLOW | FE\_UNDERFLOW) is non-  
 51505        zero, an error has occurred.

51506 **RETURN VALUE**

51507        Upon successful completion, these functions shall return the relevant Bessel value of *x* of the  
 51508        second kind.

51509        If *x* is NaN, NaN shall be returned.

51510        If the *x* argument to these functions is negative, *-HUGE\_VAL* or NaN shall be returned, and a  
 51511        domain error may occur.

51512        If *x* is 0.0, *-HUGE\_VAL* shall be returned and a range error may occur.

51513        If the correct result would cause underflow, 0.0 shall be returned and a range error may occur.

51514        If the correct result would cause overflow, *-HUGE\_VAL* or 0.0 shall be returned and a range  
 51515        error may occur.

51516 **ERRORS**

51517        These functions may fail if:

51518        Domain Error    The value of *x* is negative.

51519                        If the integer expression (*math\_errhandling* & MATH\_ERRNO) is non-zero, |  
 51520                        then *errno* shall be set to [EDOM]. If the integer expression (*math\_errhandling* |  
 51521                        & MATH\_ERREXCEPT) is non-zero, then the invalid floating-point exception |  
 51522                        shall be raised. |

51523        Range Error    The value of *x* is 0.0, or the correct result would cause overflow.

51524                        If the integer expression (*math\_errhandling* & MATH\_ERRNO) is non-zero, |  
 51525                        then *errno* shall be set to [ERANGE]. If the integer expression |  
 51526                        (*math\_errhandling* & MATH\_ERREXCEPT) is non-zero, then the overflow |  
 51527                        floating-point exception shall be raised. |

51528        Range Error    The value of *x* is too large in magnitude, or the correct result would cause  
 51529        underflow.

51530                        If the integer expression (*math\_errhandling* & MATH\_ERRNO) is non-zero, |  
 51531                        then *errno* shall be set to [ERANGE]. If the integer expression |  
 51532                        (*math\_errhandling* & MATH\_ERREXCEPT) is non-zero, then the underflow |  
 51533                        floating-point exception shall be raised. |

51534 **EXAMPLES**

51535 None.

51536 **APPLICATION USAGE**

51537 On error, the expressions (`math_errhandling & MATH_ERRNO`) and (`math_errhandling &`  
51538 `MATH_ERREXCEPT`) are independent of each other, but at least one of them must be non-zero.

51539 **RATIONALE**

51540 None.

51541 **FUTURE DIRECTIONS**

51542 None.

51543 **SEE ALSO**

51544 `feclearexcept()`, `fetetestexcept()`, `isnan()`, `j0()`, the Base Definitions volume of IEEE Std 1003.1-200x, |  
51545 Section 4.18, Treatment of Error Conditions for Mathematical Functions, <**math.h**> |

51546 **CHANGE HISTORY**

51547 First released in Issue 1. Derived from Issue 1 of the SVID.

51548 **Issue 5**

51549 The DESCRIPTION is updated to indicate how an application should check for an error. This  
51550 text was previously published in the APPLICATION USAGE section.

51551 **Issue 6**

51552 The DESCRIPTION is updated to avoid use of the term “must” for application requirements.

51553 The RETURN VALUE and ERRORS sections are reworked for alignment of the error handling |  
51554 with the ISO/IEC 9899:1999 standard.

