

# **Linux Standard Base Core Specification**

## **3.1**

## **Linux Standard Base Core Specification 3.1**

Copyright © 2004, 2005 Free Standards Group

Permission is granted to copy, distribute and/or modify this document under the terms of the GNU Free Documentation License, Version 1.1; with no Invariant Sections, with no Front-Cover Texts, and with no Back-Cover Texts. A copy of the license is included in the section entitled "GNU Free Documentation License".

Portions of the text are copyrighted by the following parties:

- The Regents of the University of California
- Free Software Foundation
- Ian F. Darwin
- Paul Vixie
- BSDI (now Wind River)
- Andrew G Morgan
- Jean-loup Gailly and Mark Adler
- Massachusetts Institute of Technology

These excerpts are being used in accordance with their respective licenses.

Linux is a trademark of Linus Torvalds.

UNIX a registered trademark of the Open Group in the United States and other countries.

LSB is a trademark of the Free Standards Group in the USA and other countries.

AMD is a trademark of Advanced Micro Devices, Inc.

Intel and Itanium are registered trademarks and Intel386 is a trademarks of Intel Corporation.

PowerPC and PowerPC Architecture are trademarks of the IBM Corporation.

OpenGL is a registered trademark of Silicon Graphics, Inc.

# Contents

<b>Foreword .....</b>	<b>vii</b>
<b>Introduction .....</b>	<b>viii</b>
<b>I Introductory Elements .....</b>	<b>9</b>
1 Scope.....	10
1.1 General.....	10
1.2 Module Specific Scope.....	10
2 References .....	11
2.1 Normative References .....	11
2.2 Informative References/Bibliography .....	12
3 Requirements .....	15
3.1 Relevant Libraries .....	15
3.2 LSB Implementation Conformance .....	15
3.3 LSB Application Conformance.....	16
4 Definitions .....	18
5 Terminology .....	19
6 Documentation Conventions .....	21
7 Relationship To ISO/IEC 9945 POSIX .....	22
8 Relationship To Other Free Standards Group Specifications.....	23
<b>II Executable And Linking Format (ELF).....</b>	<b>24</b>
9 Introduction.....	25
10 Low Level System Information.....	26
10.1 Operating System Interface .....	26
10.2 Machine Interface.....	26
11 Object Format.....	27
11.1 Object Files .....	27
11.2 Sections.....	27
11.3 Special Sections.....	30
11.4 Symbol Mapping.....	36
11.5 DWARF Extensions .....	36
11.6 Exception Frames .....	39
11.7 Symbol Versioning.....	44
11.8 ABI note tag .....	48
12 Dynamic Linking .....	49
12.1 Program Loading and Dynamic Linking.....	49
12.2 Program Header .....	49
12.3 Dynamic Entries .....	49
<b>III Base Libraries .....</b>	<b>54</b>
13 Base Libraries .....	55
13.1 Introduction .....	55
13.2 Program Interpreter .....	55
13.3 Interfaces for libc .....	55
13.4 Data Definitions for libc .....	68
13.5 Interface Definitions for libc .....	120
13.6 Interfaces for libm .....	225
13.7 Data Definitions for libm.....	228
13.8 Interface Definitions for libm .....	235
13.9 Interfaces for libpthread .....	236
13.10 Data Definitions for libpthread .....	238
13.11 Interface Definitions for libpthread .....	245

13.12 Interfaces for libgcc_s .....	245
13.13 Data Definitions for libgcc_s.....	246
13.14 Interfaces for libdl .....	249
13.15 Data Definitions for libdl .....	250
13.16 Interface Definitions for libdl .....	250
13.17 Interfaces for librt.....	253
13.18 Interfaces for libcrypt.....	254
13.19 Interfaces for libpam.....	255
13.20 Data Definitions for libpam .....	256
13.21 Interface Definitions for libpam .....	257
<b>IV Utility Libraries.....</b>	<b>270</b>
14 Utility Libraries.....	271
14.1 Introduction .....	271
14.2 Interfaces for libz.....	271
14.3 Data Definitions for libz.....	272
14.4 Interface Definitions for libz .....	274
14.5 Interfaces for libncurses.....	320
14.6 Data Definitions for libncurses.....	325
14.7 Interfaces for libutil.....	334
14.8 Interface Definitions for libutil .....	334
<b>V Commands and Utilities .....</b>	<b>340</b>
15 Commands and Utilities .....	341
15.1 Commands and Utilities .....	341
15.2 Command Behavior.....	342
<b>VI Execution Environment .....</b>	<b>403</b>
16 File System Hierarchy .....	404
16.1 /dev: Device Files.....	404
16.2 /etc: Host-specific system configuration.....	404
16.3 User Accounting Databases.....	406
16.4 Path For System Administration Utilities.....	406
17 Additional Recommendations .....	407
17.1 Recommendations for applications on ownership and permissions	407
18 Additional Behaviors .....	409
18.1 Mandatory Optional Behaviors.....	409
19 Localization .....	411
19.1 Introduction .....	411
19.2 Regular Expressions.....	411
19.3 Pattern Matching Notation .....	411
<b>VII System Initialization .....</b>	<b>413</b>
20 System Initialization.....	414
20.1 Cron Jobs .....	414
20.2 Init Script Actions.....	415
20.3 Comment Conventions for Init Scripts .....	416
20.4 Installation and Removal of Init Scripts.....	418
20.5 Run Levels.....	419
20.6 Facility Names .....	420
20.7 Script Names.....	421
20.8 Init Script Functions.....	421
<b>VIII Users &amp; Groups.....</b>	<b>424</b>
21 Users & Groups.....	425
21.1 User and Group Database.....	425

21.2 User & Group Names .....	425
21.3 User ID Ranges .....	426
21.4 Rationale.....	426
<b>IX Package Format and Installation.....</b>	<b>427</b>
22 Software Installation .....	428
22.1 Introduction .....	428
22.2 Package File Format.....	428
22.3 Package Script Restrictions .....	447
22.4 Package Tools .....	447
22.5 Package Naming.....	448
22.6 Package Dependencies .....	449
22.7 Package Architecture Considerations .....	449
<b>A Alphabetical Listing of Interfaces.....</b>	<b>450</b>
A.1 libc.....	450
A.2 libcrypt.....	462
A.3 libdl.....	463
A.4 libm.....	463
A.5 libncurses .....	468
A.6 libpam .....	472
A.7 libpthread .....	472
A.8 libert .....	474
A.9 libutil .....	474
A.10 libz .....	474
<b>B Future Directions (Informative).....</b>	<b>476</b>
B.1 Introduction .....	476
B.2 Commands And Utilities .....	477
lsbinstall.....	477
<b>C GNU Free Documentation License (Informative).....</b>	<b>481</b>
C.1 PREAMBLE .....	481
C.2 APPLICABILITY AND DEFINITIONS.....	481
C.3 VERBATIM COPYING .....	482
C.4 COPYING IN QUANTITY .....	482
C.5 MODIFICATIONS .....	483
C.6 COMBINING DOCUMENTS.....	484
C.7 COLLECTIONS OF DOCUMENTS .....	485
C.8 AGGREGATION WITH INDEPENDENT WORKS .....	485
C.9 TRANSLATION .....	485
C.10 TERMINATION .....	485
C.11 FUTURE REVISIONS OF THIS LICENSE.....	486
C.12 How to use this License for your documents.....	486

## List of Figures

11-1 Version Definition Entries .....	45
11-2 Version Definition Auxiliary Entries.....	46
11-3 Version Needed Entries .....	46
11-4 Version Needed Auxiliary Entries.....	47
12-1 Dynamic Structure.....	49

## **Foreword**

1      This is version 3.1 of the Linux Standard Base Core Specification. This specification is  
2      part of a family of specifications under the general title "Linux Standard Base".  
3      Developers of applications or implementations interested in using the LSB  
4      trademark should see the Free Standards Group Certification Policy for details.

## Introduction

The LSB defines a binary interface for application programs that are compiled and packaged for LSB-conforming implementations on many different hardware architectures. Since a binary specification shall include information specific to the computer processor architecture for which it is intended, it is not possible for a single document to specify the interface for all possible LSB-conforming implementations. Therefore, the LSB is a family of specifications, rather than a single one.

This document should be used in conjunction with the documents it references. This document enumerates the system components it includes, but descriptions of those components may be included entirely or partly in this document, partly in other documents, or entirely in other reference documents. For example, the section that describes system service routines includes a list of the system routines supported in this interface, formal declarations of the data structures they use that are visible to applications, and a pointer to the underlying referenced specification for information about the syntax and semantics of each call. Only those routines not described in standards referenced by this document, or extensions to those standards, are described in the detail. Information referenced in this way is as much a part of this document as is the information explicitly included here.

The specification carries a version number of either the form  $x.y$  or  $x.y.z$ . This version number carries the following meaning:

- The first number ( $x$ ) is the major version number. All versions with the same major version number should share binary compatibility. Any addition or deletion of a new library results in a new version number. Interfaces marked as *deprecated* may be removed from the specification at a major version change.
- The second number ( $y$ ) is the minor version number. Individual interfaces may be added if all certified implementations already had that (previously undocumented) interface. Interfaces may be marked as *deprecated* at a minor version change. Other minor changes may be permitted at the discretion of the LSB workgroup.
- The third number ( $z$ ), if present, is the editorial level. Only editorial changes should be included in such versions.

Since this specification is a descriptive Application Binary Interface, and not a source level API specification, it is not possible to make a guarantee of 100% backward compatibility between major releases. However, it is the intent that those parts of the binary interface that are visible in the source level API will remain backward compatible from version to version, except where a feature marked as "Deprecated" in one release may be removed from a future release.

Implementors are strongly encouraged to make use of symbol versioning to permit simultaneous support of applications conforming to different releases of this specification.

## I Introductory Elements

# 1 Scope

## 1.1 General

1      The Linux Standard Base (LSB) defines a system interface for compiled applications  
2      and a minimal environment for support of installation scripts. Its purpose is to  
3      enable a uniform industry standard environment for high-volume applications  
4      conforming to the LSB.

5      These specifications are composed of two basic parts: A common specification  
6      ("LSB-generic" or "generic LSB") describing those parts of the interface that remain  
7      constant across all implementations of the LSB, and an architecture-specific  
8      supplement ("LSB-arch" or "archLSB") describing the parts of the interface that vary  
9      by processor architecture. Together, the LSB-generic and the architecture-specific  
10     supplement for a single hardware architecture provide a complete interface  
11     specification for compiled application programs on systems that share a common  
12     hardware architecture.

13     The LSB-generic document shall be used in conjunction with an architecture-specific  
14     supplement. Whenever a section of the LSB-generic specification shall be  
15     supplemented by architecture-specific information, the LSB-generic document  
16     includes a reference to the architecture supplement. Architecture supplements may  
17     also contain additional information that is not referenced in the LSB-generic  
18     document.

19     The LSB contains both a set of Application Program Interfaces (APIs) and  
20     Application Binary Interfaces (ABIs). APIs may appear in the source code of portable  
21     applications, while the compiled binary of that application may use the larger set of  
22     ABIs. A conforming implementation shall provide all of the ABIs listed here. The  
23     compilation system may replace (e.g. by macro definition) certain APIs with calls to  
24     one or more of the underlying binary interfaces, and may insert calls to binary  
25     interfaces as needed.

26     The LSB is primarily a binary interface definition. Not all of the source level APIs  
27     available to applications may be contained in this specification.

## 1.2 Module Specific Scope

28     This is the Core module of the Linux Standards Base (LSB). This module provides  
29     the fundamental system interfaces, libraries, and runtime environment upon which  
30     all conforming applications and libraries depend.

31     Interfaces described in this module are mandatory except where explicitly listed  
32     otherwise. Core interfaces may be supplemented by other modules; all modules are  
33     built upon the core.

## 2 References

### 2.1 Normative References

The following referenced documents are indispensable for the application of this document. For dated references, only the edition cited applies. For undated references, the latest edition of the referenced document (including any amendments) applies.

**Note:** Where copies of a document are available on the World Wide Web, a Uniform Resource Locator (URL) is given for informative purposes only. This may point to a more recent copy of the referenced specification, or may be out of date. Reference copies of specifications at the revision level indicated may be found at the Free Standards Group's Reference Specifications (<http://refspecs.freestandards.org>) site.

Table 2-1 Normative References

Name	Title	URL
Filesystem Hierarchy Standard	Filesystem Hierarchy Standard (FHS) 2.3	<a href="http://www.pathname.com/fhs/">http://www.pathname.com/fhs/</a>
IEC 60559/IEEE 754 Floating Point	IEC 60559:1989 Binary floating-point arithmetic for microprocessor systems	<a href="http://www.ieee.org/">http://www.ieee.org/</a>
ISO C (1999)	ISO/IEC 9899: 1999, Programming Languages --C	
ISO POSIX (2003)	ISO/IEC 9945-1:2003 Information technology -- Portable Operating System Interface (POSIX) -- Part 1: Base Definitions  ISO/IEC 9945-2:2003 Information technology -- Portable Operating System Interface (POSIX) -- Part 2: System Interfaces  ISO/IEC 9945-3:2003 Information technology -- Portable Operating System Interface (POSIX) -- Part 3: Shell and Utilities  ISO/IEC 9945-4:2003 Information technology -- Portable Operating System Interface (POSIX) -- Part 4: Rationale	<a href="http://www.unix.org/version3/">http://www.unix.org/version3/</a>

## 2 References

Name	Title	URL
	Including Technical Cor. 1: 2004	
Itanium C++ ABI	Itanium C++ ABI (Revision 1.83)	<a href="http://refspecs.freestandard.org/cxxabi-1.83.html">http://refspecs.freestandard.org/cxxabi-1.83.html</a>
Large File Support	Large File Support	<a href="http://www.UNIX-systems.org/version2/whatsnew/lfs20mar.html">http://www.UNIX-systems.org/version2/whatsnew/lfs20mar.html</a>
SUSv2	CAE Specification, January 1997, System Interfaces and Headers (XSH), Issue 5 (ISBN: 1-85912-181-0, C606)	<a href="http://www.opengroup.org/publications/catalog/un.htm">http://www.opengroup.org/publications/catalog/un.htm</a>
SUSv2 Commands and Utilities	The Single UNIX® Specification(SUS) Version 2, Commands and Utilities (XCU), Issue 5 (ISBN: 1-85912-191-8, C604)	<a href="http://www.opengroup.org/publications/catalog/un.htm">http://www.opengroup.org/publications/catalog/un.htm</a>
SVID Issue 3	American Telephone and Telegraph Company, System V Interface Definition, Issue 3 ; Morristown, NJ, UNIX Press, 1989.(ISBN 0201566524)	
SVID Issue 4	System V Interface Definition,Fourth Edition	
System V ABI	System V Application Binary Interface, Edition 4.1	<a href="http://www.caldera.com/developers/devspecs/gabi41.pdf">http://www.caldera.com/developers/devspecs/gabi41.pdf</a>
System V ABI Update	System V Application Binary Interface - DRAFT - 17 December 2003	<a href="http://www.caldera.com/developers/gabi/2003-12-17/contents.html">http://www.caldera.com/developers/gabi/2003-12-17/contents.html</a>
X/Open Curses	CAE Specification, May 1996, X/Open Curses, Issue 4, Version 2 (ISBN: 1-85912-171-3, C610), plus Corrigendum U018	<a href="http://www.opengroup.org/publications/catalog/un.htm">http://www.opengroup.org/publications/catalog/un.htm</a>

11

## 2.2 Informative References/Bibliography

12

13

14

In addition, the specifications listed below provide essential background information to implementors of this specification. These references are included for information only.

12

**Table 2-2 Other References**

Name	Title	URL
DWARF Debugging Information Format, Revision 2.0.0	DWARF Debugging Information Format, Revision 2.0.0 (July 27, 1993)	<a href="http://refspecs.freestandard.org/dwarf/dwarf-2.0.0.pdf">http://refspecs.freestandard.org/dwarf/dwarf-2.0.0.pdf</a>
DWARF Debugging Information Format, Revision 3.0.0 (Draft)	DWARF Debugging Information Format, Revision 3.0.0 (Draft)	<a href="http://refspecs.freestandard.org/dwarf/">http://refspecs.freestandard.org/dwarf/</a>
ISO/IEC TR14652	ISO/IEC Technical Report 14652:2002 Specification method for cultural conventions	
ITU-T V.42	International Telecommunication Union Recommendation V.42 (2002): Error-correcting procedures for DCEs using asynchronous-to-synchronous conversion	<a href="http://www.itu.int/rec/recommendation.asp?type=folders&amp;lang=e&amp;parent=T-REC-V.42">http://www.itu.int/rec/recommendation.asp?type=folders&amp;lang=e&amp;parent=T-REC-V.42</a>
Li18nux Globalization Specification	LI18NUX 2000 Globalization Specification, Version 1.0 with Amendment 4	<a href="http://www.li18nux.org/docs/html/LI18NUX-2000-amd4.htm">http://www.li18nux.org/docs/html/LI18NUX-2000-amd4.htm</a>
Linux Allocated Device Registry	LINUX ALLOCATED DEVICES	<a href="http://www.lanana.org/docs/device-list/devices.txt">http://www.lanana.org/docs/device-list/devices.txt</a>
PAM	Open Software Foundation, Request For Comments: 86.0 , October 1995, V. Samar & R.Schemers (SunSoft)	<a href="http://www.opengroup.org/tech/rfc/mirror-rfc/rfc86.0.txt">http://www.opengroup.org/tech/rfc/mirror-rfc/rfc86.0.txt</a>
RFC 1321: The MD5 Message-Digest Algorithm	IETF RFC 1321: The MD5 Message-Digest Algorithm	<a href="http://www.ietf.org/rfc/rfc1321.txt">http://www.ietf.org/rfc/rfc1321.txt</a>
RFC 1831/1832 RPC & XDR	IETF RFC 1831 & 1832	<a href="http://www.ietf.org/">http://www.ietf.org/</a>
RFC 1833: Binding Protocols for ONC RPC Version 2	IETF RFC 1833: Binding Protocols for ONC RPC Version 2	<a href="http://www.ietf.org/rfc/rfc1833.txt">http://www.ietf.org/rfc/rfc1833.txt</a>
RFC 1950: ZLIB Compressed Data Format Specification	IETF RFC 1950: ZLIB Compressed Data Format Specification	<a href="http://www.ietf.org/rfc/rfc1950.txt">http://www.ietf.org/rfc/rfc1950.txt</a>

## 2 References

Name	Title	URL
RFC 1951: DEFLATE Compressed Data Format Specification	IETF RFC 1951: DEFLATE Compressed Data Format Specification version 1.3	<a href="http://www.ietf.org/rfc/rfc1951.txt">http://www.ietf.org/rfc/rfc1951.txt</a>
RFC 1952: GZIP File Format Specification	IETF RFC 1952: GZIP file format specification version 4.3	<a href="http://www.ietf.org/rfc/rfc1952.txt">http://www.ietf.org/rfc/rfc1952.txt</a>
RFC 2440: OpenPGP Message Format	IETF RFC 2440: OpenPGP Message Format	<a href="http://www.ietf.org/rfc/rfc2440.txt">http://www.ietf.org/rfc/rfc2440.txt</a>
RFC 2821:Simple Mail Transfer Protocol	IETF RFC 2821: Simple Mail Transfer Protocol	<a href="http://www.ietf.org/rfc/rfc2821.txt">http://www.ietf.org/rfc/rfc2821.txt</a>
RFC 2822:Internet Message Format	IETF RFC 2822: Internet Message Format	<a href="http://www.ietf.org/rfc/rfc2822.txt">http://www.ietf.org/rfc/rfc2822.txt</a>
RFC 791:Internet Protocol	IETF RFC 791: Internet Protocol Specification	<a href="http://www.ietf.org/rfc/rfc791.txt">http://www.ietf.org/rfc/rfc791.txt</a>
RPM Package Format	RPM Package Format V3.0	<a href="http://www.rpm.org/max-rpm/s1-rpm-file-format-at-rpm-file-format.html">http://www.rpm.org/max-rpm/s1-rpm-file-format-at-rpm-file-format.html</a>
zlib Manual	zlib 1.2 Manual	<a href="http://www.gzip.org/zlib/">http://www.gzip.org/zlib/</a>

16

## 3 Requirements

### 3.1 Relevant Libraries

The libraries listed in Table 3-1 shall be available on a Linux Standard Base system, with the specified runtime names. The libraries listed in Table 3-2 are architecture specific, but shall be available on all LSB conforming systems. This list may be supplemented or amended by the architecture specific supplement.

**Table 3-1 Standard Library Names**

Library	Runtime Name
libdl	libdl.so.2
libcrypt	libcrypt.so.1
libz	libz.so.1
libncurses	libncurses.so.5
libutil	libutil.so.1
libpthread	libpthread.so.0
librt	librt.so.1
libpam	libpam.so.0
libgcc_s	libgcc_s.so.1

**Table 3-2 Standard Library Names defined in the Architecture Specific Supplement**

Library	Runtime Name
libm	See archLSB
libc	See archLSB
proginterp	See archLSB

These libraries will be in an implementation-defined directory which the dynamic linker shall search by default.

### 3.2 LSB Implementation Conformance

A conforming implementation is necessarily architecture specific, and must provide the interfaces specified by both the generic LSB Core specification and its relevant architecture specific supplement.

**Rationale:** An implementation must provide *at least* the interfaces specified in these specifications. It may also provide additional interfaces.

A conforming implementation shall satisfy the following requirements:

- A processor architecture represents a family of related processors which may not have identical feature sets. The architecture specific supplement to this specification for a given target processor architecture describes a minimum

### 3 Requirements

- acceptable processor. The implementation shall provide all features of this processor, whether in hardware or through emulation transparent to the application.
- The implementation shall be capable of executing compiled applications having the format and using the system interfaces described in this document.
  - The implementation shall provide libraries containing the interfaces specified by this document, and shall provide a dynamic linking mechanism that allows these interfaces to be attached to applications at runtime. All the interfaces shall behave as specified in this document.
  - The map of virtual memory provided by the implementation shall conform to the requirements of this document.
  - The implementation's low-level behavior with respect to function call linkage, system traps, signals, and other such activities shall conform to the formats described in this document.
  - The implementation shall provide all of the mandatory interfaces in their entirety.
  - The implementation may provide one or more of the optional interfaces. Each optional interface that is provided shall be provided in its entirety. The product documentation shall state which optional interfaces are provided.
  - The implementation shall provide all files and utilities specified as part of this document in the format defined here and in other referenced documents. All commands and utilities shall behave as required by this document. The implementation shall also provide all mandatory components of an application's runtime environment that are included or referenced in this document.
  - The implementation, when provided with standard data formats and values at a named interface, shall provide the behavior defined for those values and data formats at that interface. However, a conforming implementation may consist of components which are separately packaged and/or sold. For example, a vendor of a conforming implementation might sell the hardware, operating system, and windowing system as separately packaged items.
  - The implementation may provide additional interfaces with different names. It may also provide additional behavior corresponding to data values outside the standard ranges, for standard named interfaces.

### 3.3 LSB Application Conformance

A conforming application is necessarily architecture specific, and must conform to both the generic LSB Core specification and its relevant architecture specific supplement.

A conforming application shall satisfy the following requirements:

- Its executable files shall be either shell scripts or object files in the format defined for the Object File Format system interface.
- Its object files shall participate in dynamic linking as defined in the Program Loading and Linking System interface.
- It shall employ only the instructions, traps, and other low-level facilities defined in the Low-Level System interface as being for use by applications.

- 63     • If it requires any optional interface defined in this document in order to be  
64        installed or to execute successfully, the requirement for that optional interface  
65        shall be stated in the application's documentation.
- 66     • It shall not use any interface or data format that is not required to be provided by a  
67        conforming implementation, unless:
- 68        • If such an interface or data format is supplied by another application through  
69            direct invocation of that application during execution, that application shall be  
70            in turn an LSB conforming application.
- 71        • The use of that interface or data format, as well as its source, shall be identified  
72            in the documentation of the application.
- 73     • It shall not use any values for a named interface that are reserved for vendor  
74        extensions.
- 75       A strictly conforming application shall not require or use any interface, facility, or  
76        implementation-defined extension that is not defined in this document in order to be  
77        installed or to execute successfully.

## 4 Definitions

1       For the purposes of this document, the following definitions, as specified in the  
2       *ISO/IEC Directives, Part 2, 2001, 4th Edition*, apply:

3       can

4           be able to; there is a possibility of; it is possible to

5       cannot

6           be unable to; there is no possibility of; it is not possible to

7       may

8           is permitted; is allowed; is permissible

9       need not

10           it is not required that; no...is required

11       shall

12           is to; is required to; it is required that; has to; only...is permitted; it is necessary

13       shall not

14           is not allowed [permitted] [acceptable] [permissible]; is required to be not; is  
15           required that...be not; is not to be

16       should

17           it is recommended that; ought to

18       should not

19           it is not recommended that; ought not to

## 5 Terminology

- 1           For the purposes of this document, the following terms apply:
- 2           **archLSB**
- 3           The architectural part of the LSB Specification which describes the specific parts  
4           of the interface that are platform specific. The archLSB is complementary to the  
5           gLSB.
- 6           **Binary Standard**
- 7           The total set of interfaces that are available to be used in the compiled binary  
8           code of a conforming application.
- 9           **gLSB**
- 10          The common part of the LSB Specification that describes those parts of the  
11          interface that remain constant across all hardware implementations of the LSB.
- 12          **implementation-defined**
- 13          Describes a value or behavior that is not defined by this document but is  
14          selected by an implementor. The value or behavior may vary among  
15          implementations that conform to this document. An application should not rely  
16          on the existence of the value or behavior. An application that relies on such a  
17          value or behavior cannot be assured to be portable across conforming  
18          implementations. The implementor shall document such a value or behavior so  
19          that it can be used correctly by an application.
- 20          **Shell Script**
- 21          A file that is read by an interpreter (e.g., awk). The first line of the shell script  
22          includes a reference to its interpreter binary.
- 23          **Source Standard**
- 24          The set of interfaces that are available to be used in the source code of a  
25          conforming application.
- 26          **undefined**
- 27          Describes the nature of a value or behavior not defined by this document which  
28          results from use of an invalid program construct or invalid data input. The  
29          value or behavior may vary among implementations that conform to this  
30          document. An application should not rely on the existence or validity of the  
31          value or behavior. An application that relies on any particular value or behavior  
32          cannot be assured to be portable across conforming implementations.
- 33          **unspecified**
- 34          Describes the nature of a value or behavior not specified by this document  
35          which results from use of a valid program construct or valid data input. The  
36          value or behavior may vary among implementations that conform to this  
37          document. An application should not rely on the existence or validity of the  
38          value or behavior. An application that relies on any particular value or behavior  
39          cannot be assured to be portable across conforming implementations.

## *5 Terminology*

40           Other terms and definitions used in this document shall have the same meaning as  
41           defined in Chapter 3 of the Base Definitions volume of ISO POSIX (2003).

## 6 Documentation Conventions

1       Throughout this document, the following typographic conventions are used:

2           function()  
3                 the name of a function

4           **command**

5                 the name of a command or utility

6           CONSTANT

7                 a constant value

8           *parameter*

9                 a parameter

10          variable

11                 a variable

12       Throughout this specification, several tables of interfaces are presented. Each entry  
13      in these tables has the following format:

14          name

15                 the name of the interface

16          (symver)

17                 An optional symbol version identifier, if required.

18          [refno]

19                 A reference number indexing the table of referenced specifications that follows  
20      this table.

21      For example,

22                 

forkpty(GLIBC\_2.0) [SUSv3]

23      refers to the interface named `forkpty()` with symbol version `GLIBC_2.0` that is  
24      defined in the `SUSv3` reference.

25      **Note:** Symbol versions are defined in the architecture specific supplements only.

## 7 Relationship To ISO/IEC 9945 POSIX

1 This specification includes many interfaces described in ISO POSIX (2003). Unless  
2 otherwise specified, such interfaces should behave exactly as described in that  
3 specification. Any conflict between the requirements described here and the ISO  
4 POSIX (2003) standard is unintentional, except as explicitly noted otherwise.

5 **Note:** In addition to the differences noted inline in this specification, PDTR 24715 has  
6 extracted the differences between this specification and ISO POSIX (2003) into a single  
7 place. It is the long term plan of the Free Standards Group to converge the LSB Core  
8 Specification with ISO/IEC 9945 POSIX.

9 The LSB Specification Authority is responsible for deciding the meaning of  
10 conformance to normative referenced standards in the LSB context. Problem Reports  
11 regarding underlying or referenced standards in any other context will be referred  
12 to the relevant maintenance body for that standard.

## **8 Relationship To Other Free Standards Group Specifications**

1      The LSB is the base for several other specification projects under the umbrella of the  
2      Free Standards Group (FSG). This specification is the foundation, and other  
3      specifications build on the interfaces defined here. However, beyond those  
4      specifications listed as Normative References, this specification has no dependencies  
5      on other FSG projects.

## **II Executable And Linking Format (ELF)**

## **9 Introduction**

1      Executable and Linking Format (ELF) defines the object format for compiled  
2      applications. This specification supplements the information found in System V ABI  
3      Update and is intended to document additions made since the publication of that  
4      document.

# 10 Low Level System Information

## 10.1 Operating System Interface

1 LSB-conforming applications shall assume that stack, heap and other allocated  
2 memory regions will be non-executable. The application must take steps to make  
3 them executable if needed.

## 10.2 Machine Interface

### 10.2.1 Data Representation

4 LSB-conforming applications shall use the data representation as defined in the  
5 Architecture specific ELF documents.

#### 10.2.1.1 Fundamental Types

6 In addition to the fundamental types specified in the architecture specific  
7 supplement, a 1 byte data type is defined here.  
8

9 **Table 10-1 Scalar Types**

Type	C	C++	sizeof	Alignment (bytes)	Architecture Representation
Integral	_Bool	bool	1	1	byte

10

# 11 Object Format

## 11.1 Object Files

1 LSB-conforming implementations shall support the object file Executable and  
2 Linking Format (ELF), which is defined by the following documents:

- 3   • System V ABI  
4   • System V ABI Update  
5   • this specification  
6   • an architecture specific supplement to this specification

7 Conforming implementations may also support other unspecified object file  
8 formats.

## 11.2 Sections

### 11.2.1 Introduction

9 As described in System V ABI, an ELF object file contains a number of *sections*.

### 11.2.2 Sections Types

10 The section header table is an array of `Elf32_Shdr` or `Elf64_Shdr` structures as  
11 described in System V ABI. The `sh_type` member shall be either a value from Table  
12 11-1, drawn from the System V ABI, or one of the additional values specified in  
13 Table 11-2.

14 A section header's `sh_type` member specifies the sections's semantics.

#### 11.2.2.1 ELF Section Types

16 The following section types are defined in the System V ABI and the System V ABI  
17 Update.

18 **Table 11-1 ELF Section Types**

Name	Value	Description
SHT_DYNAMIC	0x6	The section holds information for dynamic linking. Currently, an object file shall have only one dynamic section, but this restriction may be relaxed in the future. See 'Dynamic Section' in Chapter 5 for details.
SHT_DYNSYM	0xb	This section holds a minimal set of symbols adequate for dynamic linking. See also SHT_SYMTAB. Currently, an object file may have either a section of

Name	Value	Description
		SHT_SYMTAB type or a section of SHT_DYNSYM type, but not both. This restriction may be relaxed in the future.
SHT_FINI_ARRAY	0xf	This section contains an array of pointers to termination functions, as described in 'Initialization and Termination Functions' in Chapter 5. Each pointer in the array is taken as a parameterless procedure with a void return.
SHT_HASH	0x5	The section holds a symbol hash table. Currently, an object file shall have only one hash table, but this restriction may be relaxed in the future. See 'Hash Table' in the Chapter 5 for details.
SHT_INIT_ARRAY	0xe	This section contains an array of pointers to initialization functions, as described in 'Initialization and Termination Functions' in Chapter 5. Each pointer in the array is taken as a parameterless procedure with a void return.
SHT_NOBITS	0x8	A section of this type occupies no space in the file but otherwise resembles SHT_PROGBITS. Although this section contains no bytes, the sh_offset member contains the conceptual file offset.
SHT_NOTE	0x7	The section holds information that marks the file in some way. See 'Note Section' in Chapter 5 for details.

Name	Value	Description
SHT_NULL	0x0	This value marks the section header as inactive; it does not have an associated section. Other members of the section header have undefined values.
SHT_PREINIT_ARRAY	0x10	This section contains an array of pointers to functions that are invoked before all other initialization functions, as described in 'Initialization and Termination Functions' in Chapter 5. Each pointer in the array is taken as a parameterless procedure with a void return.
SHT_PROGBITS	0x1	The section holds information defined by the program, whose format and meaning are determined solely by the program.
SHT_REL	0x9	The section holds relocation entries without explicit addends, such as type Elf32_Rel for the 32-bit class of object files or type Elf64_Rel for the 64-bit class of object files. An object file may have multiple relocation sections. See "Relocation"
SHT_REL_A	0x4	The section holds relocation entries with explicit addends, such as type Elf32_Rela for the 32-bit class of object files or type Elf64_Rela for the 64-bit class of object files. An object file may have multiple relocation sections. `Relocation' b
SHT_STRTAB	0x3	The section holds a string table. An object file may have multiple string table

Name	Value	Description
		sections. See `String Table' below for details.
SHT_SYMTAB	0x2	This section holds a symbol table. Currently, an object file may have either a section of SHT_SYMTAB type or a section of SHT_DYNSYM type, but not both. This restriction may be relaxed in the future. Typically, SHT_SYMTAB provides symbols for link editing, though it may also be used for dynamic linking. As a complete symbol table, it may contain many symbols unnecessary for dynamic linking.

19

20

### 11.2.2 Additional Section Types

21

The following additional section types are defined here.

22

**Table 11-2 Additional Section Types**

Name	Value	Description
SHT_GNU_verdef	0x6fffffd	This section contains the symbol versions that are provided.
SHT_GNU_verneed	0x6fffffe	This section contains the symbol versions that are required.
SHT_GNU_versym	0x6fffffff	This section contains the Symbol Version Table.

23

## 11.3 Special Sections

24

### 11.3.1 Special Sections

25

Various sections hold program and control information. Sections in the lists below are used by the system and have the indicated types and attributes.

26

#### 11.3.1.1 ELF Special Sections

27

28

The following sections are defined in the System V ABI and the System V ABI Update.

**Table 11-3 ELF Special Sections**

Name	Type	Attributes
.bss	SHT_NOBITS	SHF_ALLOC+SHF_WRITE
.comment	SHT_PROGBITS	0
.data	SHT_PROGBITS	SHF_ALLOC+SHF_WRITE
.data1	SHT_PROGBITS	SHF_ALLOC+SHF_WRITE
.debug	SHT_PROGBITS	0
.dynamic	SHT_DYNAMIC	SHF_ALLOC+SHF_WRITE
.dynstr	SHT_STRTAB	SHF_ALLOC
.dynsym	SHT_DYNSYM	SHF_ALLOC
.fini	SHT_PROGBITS	SHF_ALLOC+SHF_EXECINSTR
.fini_array	SHT_FINI_ARRAY	SHF_ALLOC+SHF_WRITE
.hash	SHT_HASH	SHF_ALLOC
.init	SHT_PROGBITS	SHF_ALLOC+SHF_EXECINSTR
.init_array	SHT_INIT_ARRAY	SHF_ALLOC+SHF_WRITE
.interp	SHT_PROGBITS	SHF_ALLOC
.line	SHT_PROGBITS	0
.note	SHT_NOTE	0
.preinit_array	SHT_PREINIT_ARRAY	SHF_ALLOC+SHF_WRITE
.rodata	SHT_PROGBITS	SHF_ALLOC
.rodata1	SHT_PROGBITS	SHF_ALLOC
.shstrtab	SHT_STRTAB	0
.strtab	SHT_STRTAB	SHF_ALLOC
.symtab	SHT_SYMTAB	SHF_ALLOC
.tbss	SHT_NOBITS	SHF_ALLOC+SHF_WRITE+SHF_TLS
.tdata	SHT_PROGBITS	SHF_ALLOC+SHF_WRITE+SHF_TLS

Name	Type	Attributes
.text	SHT_PROGBITS	SHF_ALLOC+SHF_EXE CINSTR

- 30 .bss
- 32     This section holds data that contributes to the program's memory image. The  
 33     program may treat this data as uninitialized. However, the system shall  
 34     initialize this data with zeroes when the program begins to run. The section  
 35     occupies no file space, as indicated by the section type, SHT\_NOBITS
- 36 .comment
- 37     This section holds version control information.
- 38 .data
- 39     This section holds initialized data that contribute to the program's memory  
 40     image.
- 41 .data1
- 42     This section holds initialized data that contribute to the program's memory  
 43     image.
- 44 .debug
- 45     This section holds information for symbolic debugging. The contents are  
 46     unspecified. All section names with the prefix .debug hold information for  
 47     symbolic debugging. The contents of these sections are unspecified.
- 48 .dynamic
- 49     This section holds dynamic linking information. The section's attributes will  
 50     include the SHF\_ALLOC bit. Whether the SHF\_WRITE bit is set is processor  
 51     specific. See Chapter 5 for more information.
- 52 .dynstr
- 53     This section holds strings needed for dynamic linking, most commonly the  
 54     strings that represent the names associated with symbol table entries. See  
 55     Chapter 5 for more information.
- 56 .dynsym
- 57     This section holds the dynamic linking symbol table, as described in 'Symbol  
 58     Table'. See Chapter 5 for more information.
- 59 .fini
- 60     This section holds executable instructions that contribute to the process  
 61     termination code. That is, when a program exits normally, the system arranges  
 62     to execute the code in this section.
- 63 .fini\_array
- 64     This section holds an array of function pointers that contributes to a single  
 65     termination array for the executable or shared object containing the section.

```

66 .hash
67     This section holds a symbol hash table. See `Hash Table' in Chapter 5 for more
68     information.

69 .init
70     This section holds executable instructions that contribute to the process
71     initialization code. When a program starts to run, the system arranges to
72     execute the code in this section before calling the main program entry point
73     (called main for C programs)

74 .init_array
75     This section holds an array of function pointers that contributes to a single
76     initialization array for the executable or shared object containing the section.

77 .interp
78     This section holds the path name of a program interpreter. If the file has a
79     loadable segment that includes relocation, the sections' attributes will include
80     the SHF_ALLOC bit; otherwise, that bit will be off. See Chapter 5 for more
81     information.

82 .line
83     This section holds line number information for symbolic debugging, which
84     describes the correspondence between the source program and the machine
85     code. The contents are unspecified.

86 .note
87     This section holds information in the format that `Note Section' in Chapter 5
88     describes of the System V Application Binary Interface, Edition 4.1.

89 .preinit_array
90     This section holds an array of function pointers that contributes to a single
91     pre-initialization array for the executable or shared object containing the
92     section.

93 .rodata
94     This section holds read-only data that typically contribute to a non-writable
95     segment in the process image. See `Program Header' in Chapter 5 for more
96     information.

97 .rodata1
98     This section hold sread-only data that typically contribute to a non-writable
99     segment in the process image. See `Program Header' in Chapter 5 for more
100    information.

101 .shstrtab
102     This section holds section names.

103 .strtab
104     This section holds strings, most commonly the strings that represent the names
105     associated with symbol table entries. If the file has a loadable segment that

```

106                   includes the symbol string table, the section's attributes will include the  
 107                   SHF\_ALLOC bit; otherwi

108                 .symtab  
 109                 This section holds a symbol table, as 'Symbol Table'. in this chapter describes. If  
 110                 the file has a loadable segment that includes the symbol table, the section's  
 111                 attributes will include the SHF\_ALLOC bit; otherwise, that bit will be off.

112                 .tbss  
 113                 This section holds uninitialized thread-local data that contribute to the  
 114                 program's memory image. By definition, the system initializes the data with  
 115                 zeros when the data is instantiated for each new execution flow. The section  
 116                 occupies no file space, as indicated by the section type, SHT\_NOBITS.  
 117                 Implementations need not support thread-local storage.

118                 .tdata  
 119                 This section holds initialized thread-local data that contributes to the program's  
 120                 memory image. A copy of its contents is instantiated by the system for each new  
 121                 execution flow. Implementations need not support thread-local storage.

122                 .text  
 123                 This section holds the 'text,' or executable instructions, of a program.

### 11.3.1.2 Additional Special Sections

125                 Object files in an LSB conforming application may also contain one or more of the  
 126                 additional special sections described below.

127                 **Table 11-4 Additional Special Sections**

Name	Type	Attributes
.ctors	SHT_PROGBITS	SHF_ALLOC+SHF_WRITE
.data.rel.ro	SHT_PROGBITS	SHF_ALLOC+SHF_WRITE
.dtors	SHT_PROGBITS	SHF_ALLOC+SHF_WRITE
.eh_frame	SHT_PROGBITS	SHF_ALLOC
.eh_frame_hdr	SHT_PROGBITS	SHF_ALLOC
.gcc_except_table	SHT_PROGBITS	SHF_ALLOC
.gnu.version	SHT_GNU_versym	SHF_ALLOC
.gnu.version_d	SHT_GNU_verdef	SHF_ALLOC
.gnu.version_r	SHT_GNU_verneed	SHF_ALLOC
.got.plt	SHT_PROGBITS	SHF_ALLOC+SHF_WRITE
.jcr	SHT_PROGBITS	SHF_ALLOC+SHF_WRITE

Name	Type	Attributes
.note.ABI-tag	SHT_NOTE	SHF_ALLOC
.stab	SHT_PROGBITS	0
.stabstr	SHT_STRTAB	0

128

129 .ctors

130

This section contains a list of global constructor function pointers.

131

.data.rel.ro

132

This section holds initialized data that contribute to the program's memory image. This section may be made read-only after relocations have been applied.

133

.dtors

134

This section contains a list of global destructor function pointers.

135

.eh\_frame

136

This section contains information necessary for frame unwinding during exception handling. See Section 11.6.1.

137

.eh\_frame\_hdr

138

This section contains a pointer to the .eh\_frame section which is accessible to the runtime support code of a C++ application. This section may also contain a binary search table which may be used by the runtime support code to more efficiently access records in the .eh\_frame section. See Section 11.6.2.

139

.gcc\_except\_table

140

This section holds Language Specific Data.

141

.gnu.version

142

This section contains the Symbol Version Table. See Section 11.7.2.

143

.gnu.version\_d

144

This section contains the Version Definitions. See Section 11.7.3.

145

.gnu.version\_r

146

This section contains the Version Requirements. See Section 11.7.4.

147

.got.plt

148

This section holds the read-only portion of the Global Offset Table. This section may be made read-only after relocations have been applied.

149

.jcr

150

This section contains information necessary for registering compiled Java classes. The contents are compiler-specific and used by compiler initialization functions.

151

.note.ABI-tag

152

Specify ABI details. See Section 11.8.

161            .stab  
 162            This section contains debugging information. The contents are not specified as  
 163            part of the LSB.  
 164            .stabstr  
 165            This section contains strings associated with the debugging infomation  
 166            contained in the .stab section.

## 11.4 Symbol Mapping

### 11.4.1 Introduction

167            Symbols in a source program are translated by the compilation system into symbols  
 168            that exist in the object file.

#### 11.4.1.1 C Language

170            External C symbols shall be unchanged in an object file's symbol table.

## 11.5 DWARF Extensions

171            The LSB does not specify debugging information, however, some additional sections  
 172            contain information which is encoded using the the encoding as specified by  
 173            DWARF Debugging Information Format, Revision 2.0.0 with extensions defined  
 174            here.

175            **Note:** The extensions specified here also exist in DWARF Debugging Information  
 176            Format, Revision 3.0.0 (Draft). It is expected that future versions of the LSB will reference  
 177            the final version of that document, and that the definitions here will be taken from that  
 178            document instead of being specified here.

### 11.5.1 DWARF Exception Header Encoding

179            The DWARF Exception Header Encoding is used to describe the type of data used in  
 180            the .eh\_frame and .eh\_frame\_hdr section. The upper 4 bits indicate how the value  
 181            is to be applied. The lower 4 bits indicate the format of the data.

182            **Table 11-5 DWARF Exception Header value format**

Name	Value	Meaning
DW_EH_PE_absptr	0x00	The Value is a literal pointer whose size is determined by the architecture.
DW_EH_PE_uleb128	0x01	Unsigned value is encoded using the Little Endian Base 128 (LEB128) as defined by DWARF Debugging Information Format, Revision 2.0.0.
DW_EH_PE_udata2	0x02	A 2 bytes unsigned value.

183

Name	Value	Meaning
DW_EH_PE_udata4	0x03	A 4 bytes unsigned value.
DW_EH_PE_udata8	0x04	An 8 bytes unsigned value.
DW_EH_PE_sleb128	0x09	Signed value is encoded using the Little Endian Base 128 (LEB128) as defined by DWARF Debugging Information Format, Revision 2.0.0.
DW_EH_PE_sdata2	0x0A	A 2 bytes signed value.
DW_EH_PE_sdata4	0x0B	A 4 bytes signed value.
DW_EH_PE_sdata8	0x0C	An 8 bytes signed value.

184

**Table 11-6 DWARF Exception Header application**

185

Name	Value	Meaning
DW_EH_PE_pcrel	0x10	Value is relative to the current program counter.
DW_EH_PE_textrel	0x20	Value is relative to the beginning of the .text section.
DW_EH_PE_datarel	0x30	Value is relative to the beginning of the .got or .eh_frame_hdr section.
DW_EH_PE_funcrel	0x40	Value is relative to the beginning of the function.
DW_EH_PE_aligned	0x50	Value is aligned to an address unit sized boundary.

186

187

One special encoding, 0xff (DW\_EH\_PE\_omit), shall be used to indicate that no value is present.

### 11.5.2 DWARF CFI Extensions

188

189

190

In addition to the Call Frame Instructions defined in section 6.4.2 of DWARF Debugging Information Format, Revision 2.0.0, the following additional Call Frame Instructions may also be used.

191

**Table 11-7 Additional DWARF Call Frame Instructions**

Name	Value	Meaning
DW_CFA_expression	0x10	The DW_CFA_expression

Name	Value	Meaning
		instruction takes two operands: an unsigned LEB128 value representing a register number, and a DW_FORM_block value representing a DWARF expression. The required action is to establish the DWARF expression as the means by which the address in which the given register contents are found may be computed. The value of the CFA is pushed on the DWARF evaluation stack prior to execution of the DWARF expression. The DW_OP_call2, DW_OP_call4, DW_OP_call_ref and DW_OP_push_object_address DWARF operators (see Section 2.4.1 of DWARF Debugging Information Format, Revision 2.0.0) cannot be used in such a DWARF expression.
DW_CFA_offset_extended_sf	0x11	The DW_CFA_offset_extended_sf instruction takes two operands: an unsigned LEB128 value representing a register number and a signed LEB128 factored offset. This instruction is identical to DW_CFA_offset_extended except that the second operand is signed.
DW_CFA_def_cfa_sf	0x12	The DW_CFA_def_cfa_sf instruction takes two operands: an unsigned LEB128 value representing a register number and a signed LEB128 factored offset.

Name	Value	Meaning
		This instruction is identical to DW_CFA_def_cfa except that the second operand is signed and factored.
DW_CFA_def_cfa_offset_sf	0x13	The DW_CFA_def_cfa_offset_sf instruction takes a signed LEB128 operand representing a factored offset. This instruction is identical to DW_CFA_def_cfa_offset except that the operand is signed and factored.
DW_CFA_GNU_args_size	0x2e	The DW_CFA_GNU_args_size instruction takes an unsigned LEB128 operand representing an argument size. This instruction specifies the total of the size of the arguments which have been pushed onto the stack.
DW_CFA_GNU_negative_offset_extended	0x2f	The DW_CFA_def_cfa_sf instruction takes two operands: an unsigned LEB128 value representing a register number and an unsigned LEB128 which represents the magnitude of the offset. This instruction is identical to DW_CFA_offset_extended_sf except that the operand is subtracted to produce the offset. This instruction is obsoleted by DW_CFA_offset_extended_sf.

192

## 11.6 Exception Frames

193  
194

When using languages that support exceptions, such as C++, additional information must be provided to the runtime environment that describes the call frames that

195 must be unwound during the processing of an exception. This information is  
 196 contained in the special sections `.eh_frame` and `.eh_framehdr`.

197 **Note:** The format of the `.eh_frame` section is similar in format and purpose to  
 198 the `.debug_frame` section which is specified in DWARF Debugging Information Format,  
 199 Revision 3.0.0 (Draft). Readers are advised that there are some subtle difference, and care  
 200 should be taken when comparing the two sections.

### 11.6.1 The `.eh_frame` section

201 The `.eh_frame` section shall contain 1 or more Call Frame Information (CFI) records.  
 202 The number of records present shall be determined by size of the section as  
 203 contained in the section header. Each CFI record contains a Common Information  
 204 Entry (CIE) record followed by 1 or more Frame Description Entry (FDE) records.  
 205 Both CIEs and FDEs shall be aligned to an addressing unit sized boundary.

206 **Table 11-8 Call Frame Information Format**

Common Information Entry Record
Frame Description Entry Record(s)

208 **11.6.1.1 The Common Information Entry Format**

209 **Table 11-9 Common Information Entry Format**

Length	Required
Extended Length	Optional
CIE ID	Required
Version	Required
Augmentation String	Required
Code Alignment Factor	Required
Data Alignment Factor	Required
Return Address Register	Required
Augmentation Data Length	Optional
Augmentation Data	Optional
Initial Instructions	Required
Padding	

211 *Length*

212 A 4 byte unsigned value indicating the length in bytes of the CIE structure, not  
 213 including the *Length* field itself. If *Length* contains the value 0xffffffff, then the  
 214 length is contained in the *Extended Length* field. If *Length* contains the value 0,  
 215 then this CIE shall be considered a terminator and processing shall end.

216 *Extended Length*

217 A 8 byte unsigned value indicating the length in bytes of the CIE structure, not  
 218 including the *Length* and *Extended Length* fields.

219           *CIE ID*  
 220           A 4 byte unsigned value that is used to distinguish CIE records from FDE  
 221           records. This value shall always be 0, which indicates this record is a CIE.

222           *Version*  
 223           A 1 byte value that identifies the version number of the frame information  
 224           structure. This value shall be 1.

225           *Augmentation String*  
 226           This value is a NUL terminated string that identifies the augmentation to the  
 227           CIE or to the FDEs associated with this CIE. A zero length string indicates that  
 228           no augmentation data is present. The augmentation string is case sensitive and  
 229           shall be interpreted as described below.

230           *Code Alignment Factor*  
 231           An unsigned LEB128 encoded value that is factored out of all advance location  
 232           instructions that are associated with this CIE or its FDEs. This value shall be  
 233           multiplied by the delta argument of an adavance location instruction to obtain  
 234           the new location value.

235           *Data Alignment Factor*  
 236           A signed LEB128 encoded value that is factored out of all offset instructions that  
 237           are associated with this CIE or its FDEs. This value shall be multiplied by the  
 238           register offset argument of an offset instruction to obtain the new offset value.

239           *Augmentation Length*  
 240           An unsigned LEB128 encoded value indicating the length in bytes of the  
 241           Augmentation Data. This field is only present if the Augmentation String  
 242           contains the character 'z'.

243           *Augmentation Data*  
 244           A block of data whose contents are defined by the contents of the Augmentation  
 245           String as described below. This field is only present if the Augmentation String  
 246           contains the character 'z'. The size of this data is given by the Augentation  
 247           Length.

248           *Initial Instructions*  
 249           Initial set of Call Frame Instructions. The number of instructions is determined  
 250           by the remaining space in the CIE record.

251           *Padding*  
 252           Extra bytes to align the CIE structure to an addressing unit size boundary.

253           **11.6.1.1 Augmentation String Format**  
 254           The Agumentation String indicates the presence of some optional fields, and how  
 255           those fields should be intrepreted. This string is case sensitive. Each character in the  
 256           augmentation string in the CIE can be interpreted as below:

257           'z'  
 258           A 'z' may be present as the first character of the string. If present, the  
 259           Augmentation Data field shall be present. The contents of the Augmentation

260 Data shall be interpreted according to other characters in the Augmentation  
 261 String.

262 'L'

263 A 'L' may be present at any position after the first character of the string. This  
 264 character may only be present if 'z' is the first character of the string. If present,  
 265 it indicates the presence of one argument in the Augmentation Data of the CIE,  
 266 and a corresponding argument in the Augmentation Data of the FDE. The  
 267 argument in the Augmentation Data of the CIE is 1-byte and represents the  
 268 pointer encoding used for the argument in the Augmentation Data of the FDE,  
 269 which is the address of a language-specific data area (LSDA). The size of the  
 270 LSDA pointer is specified by the pointer encoding used.

271 'P'

272 A 'P' may be present at any position after the first character of the string. This  
 273 character may only be present if 'z' is the first character of the string. If present,  
 274 it indicates the presence of two arguments in the Augmentation Data of the CIE.  
 275 The first argument is 1-byte and represents the pointer encoding used for the  
 276 second argument, which is the address of a *personality routine* handler. The  
 277 personality routine is used to handle language and vendor-specific tasks. The  
 278 system unwind library interface accesses the language-specific exception  
 279 handling semantics via the pointer to the personality routine. The personality  
 280 routine does not have an ABI-specific name. The size of the personality routine  
 281 pointer is specified by the pointer encoding used.

282 'R'

283 A 'R' may be present at any position after the first character of the string. This  
 284 character may only be present if 'z' is the first character of the string. If present,  
 285 The Augmentation Data shall include a 1 byte argument that represents the  
 286 pointer encoding for the address pointers used in the FDE.

### 287 11.6.1.2 The Frame Description Entry Format

288 **Table 11-10 Frame Description Entry Format**

Length	Required
Extended Length	Optional
CIE Pointer	Required
PC Begin	Required
PC Range	Required
Augmentation Data Length	Optional
Augmentation Data	Optional
Call Frame Instructions	Required
Padding	

289 *Length*

290 A 4 byte unsigned value indicating the length in bytes of the CIE structure, not  
 291 including the *Length* field itself. If *Length* contains the value 0xffffffff, then the

293 length is contained the *Extended Length* field. If *Length* contains the value 0,  
 294 then this CIE shall be considered a terminator and processing shall end.

295 *Extended Length*

296 A 8 byte unsigned value indicating the length in bytes of the CIE structure, not  
 297 including the *Length* field itself.

298 *CIE Pointer*

299 A 4 byte unsigned value that when subtracted from the offset of the current FDE  
 300 yields the offset of the start of the associated CIE. This value shall never be 0.

301 *PC Begin*

302 An encoded value that indicates the address of the initial location associated  
 303 with this FDE. The encoding format is specified in the Augmentation Data.

304 *PC Range*

305 An absolute value that indicates the number of bytes of instructions associated  
 306 with this FDE.

307 *Augmentation Length*

308 An unsigned LEB128 encoded value indicating the length in bytes of the  
 309 Augmentation Data. This field is only present if the Augmentation String in the  
 310 associated CIE contains the character 'z'.

311 *Augmentation Data*

312 A block of data whose contents are defined by the contents of the Augmentation  
 313 String in the associated CIE as described above. This field is only present if the  
 314 Augmentation String in the associated CIE contains the character 'z'. The size of  
 315 this data is given by the Augmentation Length.

316 *Call Frame Instructions*

317 A set of Call Frame Instructions.

318 *Padding*

319 Extra bytes to align the FDE structure to an addressing unit size boundary.

### 11.6.2 The .eh\_frame\_hdr section

320 The .eh\_frame\_hdr section contains additional information about the .eh\_frame  
 321 section. A pointer to the start of the .eh\_frame data, and optionally, a binary search  
 322 table of pointers to the .eh\_frame records are found in this section.

323 Data in this section is encoded according to Section 11.5.1.

324 **Table 11-11 .eh\_frame\_hdr Section Format**

<b>Encoding</b>	<b>Field</b>
unsigned byte	version
unsigned byte	eh_frame_ptr_enc
unsigned byte	fde_count_enc
unsigned byte	table_enc

<b>Encoding</b>	<b>Field</b>
encoded	eh_frame_ptr
encoded	fde_count
	binary search table

325

version

327

Version of the .eh\_frame\_hdr format. This value shall be 1.

328

eh\_frame\_ptr\_enc

329

The encoding format of the eh\_frame\_ptr field.

330

fde\_count\_enc

331

The encoding format of the fde\_count field. A value of DW\_EH\_PE\_omit indicates the binary search table is not present.

332

table\_enc

334

The encoding format of the entries in the binary search table. A value of DW\_EH\_PE\_omit indicates the binary search table is not present.

335

336

eh\_frame\_ptr

337

The encoded value of the pointer to the start of the .eh\_frame section.

338

fde\_count

339

The encoded value of the count of entries in the binary search table.

340

binary search table

341

A binary search table containing fde\_count entries. Each entry of the table consist of two encoded values, the initial location, and the address. The entries are sorted in an increasing order by the initial location value.

342

343

## 11.7 Symbol Versioning

### 11.7.1 Introduction

344

This chapter describes the Symbol Versioning mechanism. All ELF objects may provide or depend on versioned symbols. Symbol Versioning is implemented by 3 section types: SHT\_GNU\_versym, SHT\_GNU\_verdef, and SHT\_GNU\_verneed.

345

346

The prefix Elfxx in the following descriptions and code fragments stands for either "Elf32" or "Elf64", depending on the architecture.

347

348

Versions are described by strings. The structures that are used for symbol versions also contain a member that holds the ELF hashing values of the strings. This allows for more efficient processing.

349

350

351

### 11.7.2 Symbol Version Table

352

353

354

The special section .gnu.version which has a section type of SHT\_GNU\_versym shall contain the Symbol Version Table. This section shall have the same number of entries as the Dynamic Symbol Table in the .dynsym section.

355       The .gnu.version section shall contain an array of elements of type Elfxx\_Half.  
 356       Each entry specifies the version defined for or required by the corresponding symbol  
 357       in the Dynamic Symbol Table.

358       The values in the Symbol Version Table are specific to the object in which they are  
 359       located. These values are identifiers that are provided by the the vna\_other member  
 360       of the Elfxx\_Vernaux structure or the vd\_ndx member of the Elfxx\_Verdef  
 361       structure.

362       The values 0 and 1 are reserved.

363       0

364           The symbol is local, not available outside the object.

365       1

366           The symbol is defined in this object and is globally available.

367       All other values are used to identify version strings located in one of the other  
 368       Symbol Version sections. The value itself is not the version associated with the  
 369       symbol. The string identified by the value defines the version of the symbol.

### 11.7.3 Version Definitions

370       The special section .gnu.version\_d which has a section type of SHT\_GNU\_verdef  
 371       shall contain symbol version definitions. The number of entries in this section shall  
 372       be contained in the DT\_VERDEFNUM entry of the Dynamic Section .dynamic. The  
 373       sh\_link member of the section header (see figure 4-8 in the System V ABI) shall  
 374       point to the section that contains the strings referenced by this section.

375       The section shall contain an array of Elfxx\_Verdef structures, as described in  
 376       Figure 11-1, optionally followed by an array of Elfxx\_Verdaux structures, as  
 377       defined in Figure 11-2.

```
378       typedef struct {
379           Elfxx_Half     vd_version;
380           Elfxx_Half     vd_flags;
381           Elfxx_Half     vd_ndx;
382           Elfxx_Half     vd_cnt;
383           Elfxx_Word    vd_hash;
384           Elfxx_Word    vd_aux;
385           Elfxx_Word    vd_next;
386       } Elfxx_Verdef;
```

387       **Figure 11-1 Version Definition Entries**

388       *vd\_version*

389           Version revision. This field shall be set to 1.

390       *vd\_flags*

391           Version information flag bitmask.

392       *vd\_ndx*

393           Version index numeric value referencing the SHT\_GNU\_versym section.

394       *vd\_cnt*

395           Number of associated verdaux array entries.

```

396      vd_hash
397          Version name hash value (ELF hash function).
398
399      vd_aux
400          Offset in bytes to a corresponding entry in an array of Elfxx_Verdaux
401          structures as defined in Figure 11-2
402
403      vd_next
404          Offset to the next verdef entry, in bytes.
405
406      typedef struct {
407          Elfxx_Word    vda_name;
408          Elfxx_Word    vda_next;
409      } Elfxx_Verdaux;

```

**Figure 11-2 Version Definition Auxiliary Entries**

```

407
408      vda_name
409          Offset to the version or dependency name string in the section header, in bytes.
410
411      vda_next
412          Offset to the next verdaux entry, in bytes.

```

**11.7.4 Version Requirements**

The special section .gnu.version\_r which has a section type of SHT\_GNU\_verneed shall contain required symbol version definitions. The number of entries in this section shall be contained in the DT\_VERNEEDNUM entry of the Dynamic Section .dynamic. The *sh\_link* member of the section header (see figure 4-8 in System V ABI) shall point to the section that contains the strings referenced by this section.

The section shall contain an array of Elfxx\_Verneed structures, as described in Figure 11-3, optionally followed by an array of Elfxx\_Vernaux structures, as defined in Figure 11-4.

```

421      typedef struct {
422          Elfxx_Half    vn_version;
423          Elfxx_Half    vn_cnt;
424          Elfxx_Word    vn_file;
425          Elfxx_Word    vn_aux;
426          Elfxx_Word    vn_next;
427      } Elfxx_Verneed;

```

**Figure 11-3 Version Needed Entries**

```

428
429      vn_version
430          Version of structure. This value is currently set to 1, and will be reset if the
431          versioning implementation is incompatibly altered.
432
433      vn_cnt
434          Number of associated verneed array entries.
435
436      vn_file
437          Offset to the file name string in the section header, in bytes.

```

```

436     vn_aux
437         Offset to a corresponding entry in the vernaux array, in bytes.

438     vn_next
439         Offset to the next verneed entry, in bytes.

440     typedef struct {
441         Elfxx_Word    vna_hash;
442         Elfxx_Half   vna_flags;
443         Elfxx_Half   vna_other;
444         Elfxx_Word   vna_name;
445         Elfxx_Word   vna_next;
446     } Elfxx_Vernaux;

```

**Figure 11-4 Version Needed Auxiliary Entries**

```

447     vna_hash
448         Dependency name hash value (ELF hash function).

449     vna_flags
450         Dependency information flag bitmask.

451     vna_other
452         Object file version identifier used in the .gnu.version symbol version array. Bit
453             number 15 controls whether or not the object is hidden; if this bit is set, the
454             object cannot be used and the static linker will ignore the symbol's presence in
455             the object.

456     vna_name
457         Offset to the dependency name string in the section header, in bytes.

458     vna_next
459         Offset to the next vernaux entry, in bytes.

```

### 11.7.5 Startup Sequence

When loading a sharable object the system shall analyze version definition data from the loaded object to assure that it meets the version requirements of the calling object. This step is referred to as definition testing. The dynamic loader shall retrieve the entries in the caller's `Elfxx_Verneed` array and attempt to find matching definition information in the loaded `Elfxx_Verdef` table.

Each object and dependency shall be tested in turn. If a symbol definition is missing and the `vna_flags` bit for `VER_FLG_WEAK` is not set, the loader shall return an error and exit. If the `vna_flags` bit for `VER_FLG_WEAK` is set in the `Elfxx_Vernaux` entry, and the loader shall issue a warning and continue operation.

When the versions referenced by undefined symbols in the loaded object are found, version availability is certified. The test completes without error and the object shall be made available.

### 11.7.6 Symbol Resolution

When symbol versioning is used in an object, relocations extend definition testing beyond the simple match of symbol name strings: the version of the reference shall also equal the name of the definition.

476       The same index that is used in the symbol table can be referenced in the  
 477       SHT\_GNU\_versym section, and the value of this index is then used to acquire name  
 478       data. The corresponding requirement string is retrieved from the Elfxx\_Verneed  
 479       array, and likewise, the corresponding definition string from the Elfxx\_Verdef  
 480       table.

481       If the high order bit (bit number 15) of the version symbol is set, the object cannot be  
 482       used and the static linker shall ignore the symbol's presence in the object.

483       When an object with a reference and an object with the definition are being linked,  
 484       the following rules shall govern the result:

- 485       • The object with the reference and the object with the definitions both use  
         486       versioning. All described matching is processed in this case. A fatal error shall be  
         487       triggered when no matching definition can be found in the object whose name is  
         488       the one referenced by the *vn\_name* element in the Elfxx\_Verneed entry.
- 489       • The object with the reference does not use versioning, while the object with the  
         490       definitions does. In this instance, only the definitions with index numbers 1 and 2  
         491       will be used in the reference match, the same identified by the static linker as the  
         492       base definition. In cases where the static linker was not used, such as in calls to  
         493       dlopen(), a version that does not have the base definition index shall be  
         494       acceptable if it is the only version for which the symbol is defined.
- 495       • The object with the reference uses versioning, but the object with the definitions  
         496       specifies none. A matching symbol shall be accepted in this case. A fatal error shall  
         497       be triggered if a corruption in the required symbols list obscures an outdated  
         498       object file and causes a match on the object filename in the Elfxx\_Verneed entry.
- 499       • Neither the object with the reference nor the object with the definitions use  
         500       versioning. The behavior in this instance shall default to pre-existing symbol rules.

## 11.8 ABI note tag

501       Every executable shall contain a section named .note.ABI-tag of type SHT\_NOTE.  
 502       This section is structured as a note section as documented in the ELF spec. The  
 503       section shall contain at least the following entry. The name field (namesz/name)  
 504       contains the string "GNU". The type field shall be 1. The descsz field shall be at least  
 505       16, and the first 16 bytes of the desc field shall be as follows.

506       The first 32-bit word of the desc field shall be 0 (this signifies a Linux executable).  
 507       The second, third, and fourth 32-bit words of the desc field contain the earliest  
 508       compatible kernel version. For example, if the 3 words are 2, 2, and 5, this signifies a  
 509       2.2.5 kernel.

## 12 Dynamic Linking

### 12.1 Program Loading and Dynamic Linking

1 LSB-conforming implementations shall support the object file information and  
2 system actions that create running programs as specified in the System V ABI and  
3 System V ABI Update and as further required by this specification and its  
4 architecture specific supplement.

5 Any shared object that is loaded shall contain sufficient DT\_NEEDED records to  
6 satisfy the symbols on the shared library.

### 12.2 Program Header

7 In addition to the Segment Types defined in the System V ABI and System V ABI  
8 Update the following Segment Types shall also be supported.

9 **Table 12-1 Linux Segment Types**

Name	Value
PT_GNU_EH_FRAME	0x6474e550
PT_GNU_STACK	0x6474e551
PT_GNU_RELRO	0x6474e552

10 PT\_GNU\_EH\_FRAME

11 The array element specifies the location and size of the exception handling  
12 information as defined by the .eh\_frame\_hdr section.

13 PT\_GNU\_STACK

14 The *p\_flags* member specifies the permissions on the segment containing the  
15 stack and is used to indicate whether the stack should be executable. The absence  
16 of this header indicates that the stack will be executable.

17 PT\_GNU\_RELRO

18 The array element specifies the location and size of a segment which may be  
19 made read-only after relocation have been processed.

### 12.3 Dynamic Entries

#### 12.3.1 Introduction

20 As described in System V ABI, if an object file participates in dynamic linking,  
21 its program header table shall have an element of type PT\_DYNAMIC. This 'segment'  
22 contains the .dynamic section. A special symbol, \_DYNAMIC, labels the section, which  
23 contains an array of the following structures.

```
24
25     typedef struct {
26         Elf32_Sword      d_tag;
27         union {
28             Elf32_Word      d_val;
29             Elf32_Addr      d_ptr;
30         } d_un;
31     } Elf32_Dyn;
```

```
32      extern Elf32_Dyn      _DYNAMIC[ ];
33
34      typedef struct {
35          Elf64_Sxword    d_tag;
36          union {
37              Elf64_Xword   d_val;
38              Elf64_Addr    d_ptr;
39          } d_un;
40      } Elf64_Dyn;
41
42      extern Elf64_Dyn      _DYNAMIC[ ];
43
44 Figure 12-1 Dynamic Structure
```

For each object with this type, *d\_tag* controls the interpretation of *d\_un*.

## 12.3.2 Dynamic Entries

### 12.3.2.1 ELF Dynamic Entries

The following dynamic entries are defined in the System V ABI and System V ABI Update.

```
49      DT_BIND_NOW
50          Process relocations of object
51
52      DT_DEBUG
53          For debugging; unspecified
54
55      DT_FINI
56          Address of termination function
57
58      DT_HASH
59          Address of symbol hash table
60
61      DT_HIPROC
62          End of processor-specific
63
64      DT_INIT
65          Address of init function
66
67      DT_JMPREL
68          Address of PLT relocs
69
70      DT_LOPROC
71          Start of processor-specific
72
73      DT_NEEDED
74          Name of needed library
75
76      DT_NULL
77          Marks end of dynamic section
```

```

69      DT_PLTREL
70          Type of reloc in PLT
71      DT_PLTRELSZ
72          Size in bytes of PLT relocs
73      DT_REL
74          Address of Rel relocs
75      DT_RELA
76          Address of Rela relocs
77      DT_RELENT
78          Size of one Rela reloc
79      DT_RELASZ
80          Total size of Rela relocs
81      DT_RELENT
82          Size of one Rel reloc
83      DT_RELDSZ
84          Total size of Rel relocs
85      DT_RPATH
86          Library search path
87      DT SONAME
88          Name of shared object
89      DT_STRSZ
90          Size of string table
91      DT_STRTAB
92          Address of string table
93      DT_SYMBOLIC
94          Start symbol search here
95      DT_SYMENT
96          Size of one symbol table entry
97      DT_SYMTAB
98          Address of symbol table
99      DT_TEXTREL
100         Reloc might modify .text

```

101           **12.3.2.2 Additional Dynamic Entries**

102           An LSB conforming object may also use the following additional Dynamic Entry  
103           types.

104           DTADDRNGHI

105           Values from DTADDRNGLO through DTADDRNGHI are reserved for  
106           definition by an archLSB.

107           DTADDRNGLO

108           Values from DTADDRNGLO through DTADDRNGHI are reserved for  
109           definition by an archLSB.

110           DT\_AUXILIARY

111           Shared object to load before self

112           DT\_FILTER

113           Shared object to get values from

114           DT\_INIT\_ARRAY

115           The address of an array of pointers to termination functions.

116           DT\_INIT\_ARRAYSZ

117           Size in bytes of DT\_INIT\_ARRAY

118           DT\_HIOS

119           Values from DT\_LOOS through DT\_HIOS are reserved for definition by specific  
120           operating systems.

121           DT\_INIT\_ARRAY

122           The address of an array of pointers to initialization functions.

123           DT\_INIT\_ARRAYSZ

124           Size in bytes of DT\_INIT\_ARRAY

125           DT\_LOOS

126           Values from DT\_LOOS through DT\_HIOS are reserved for definition by specific  
127           operating systems.

128           DT\_NUM

129           Number of dynamic entry tags defined (excepting reserved ranges).

130           DT\_POSFLAG\_1

131           Flags for DT\_\* entries, effecting the following DT\_\* entry

132           DT\_RELCOUNT

133           All Elf32\_Rel R\_\*\_RELATIVE relocations have been placed into a single block  
134           and this entry specifies the number of entries in that block. This permits ld.so.1  
135           to streamline the processing of RELATIVE relocations.

136	DT_RUNPATH	
137		null-terminated library search path string
138	DT_SYMINENT	
139		Entry size of syminfo
140	DT_SYMINFO	
141		Address of the Syminfo table.
142	DT_SYMINSZ	
143		Size of syminfo table (in bytes)
144	DT_VALRNGHI	
145		Entries which fall between DT_VALRNGHI & DT_VALRNGLO use the
146		Dyn.d_un.d_val field of the Elf*_Dyn structure.
147	DT_VALRNGLO	
148		Entries which fall between DT_VALRNGHI & DT_VALRNGLO use the
149		Dyn.d_un.d_val field of the Elf*_Dyn structure.
150	DT_VERDEF	
151		Address of version definition table
152	DT_VERDEFNUM	
153		Number of version definitions
154	DT_VERNEED	
155		Address of table with needed versions
156	DT_VERNEEDNUM	
157		Number of needed versions
158	DT_VERSYM	
159		Address of the table provided by the .gnu.version section.

### **III Base Libraries**

## 13 Base Libraries

### 13.1 Introduction

1 An LSB-conforming implementation shall support the following base libraries  
2 which provide interfaces for accessing the operating system, processor and other  
3 hardware in the system.

- 4     • libc  
5     • libm  
6     • libgcc\_s  
7     • libdl  
8     • librt  
9     • libcrypt  
10    • libpam

11 There are three main parts to the definition of each of these libraries.

12 The "Interfaces" section defines the required library name and version, and the  
13 required public symbols (interfaces and global data), as well as symbol versions, if  
14 any.

15 The "Interface Definitions" section provides complete or partial definitions of certain  
16 interfaces where either this specification is the source specification, or where there  
17 are variations from the source specification. If an interface definition requires one or  
18 more header files, one of those headers shall include the function prototype for the  
19 interface.

20 For source definitions of interfaces which include a reference to a header file, the  
21 contents of such header files form a part of the specification. The "Data Definitions"  
22 section provides the binary-level details for the header files from the source  
23 specifications, such as values for macros and enumerated types, as well as structure  
24 layouts, sizes and padding, etc. These data definitions, although presented in the  
25 form of header files for convenience, should not be taken as representing complete  
26 header files, as they are a supplement to the source specifications. Application  
27 developers should follow the guidelines of the source specifications when  
28 determining which header files need to be included to completely resolve all  
29 references.

30 **Note:** While the Data Definitions supplement the source specifications, this specification  
31 itself does not require conforming implementations to supply any header files.

### 13.2 Program Interpreter

32 The Program Interpreter is specified in the appropriate architecture specific  
33 supplement.

### 13.3 Interfaces for libc

34 Table 13-1 defines the library name and shared object name for the libc library

35 **Table 13-1 libc Definition**

Library:	libc
----------	------

36

SONAME:	See archLSB.
---------	--------------

37

38

The behavior of the interfaces in this library is specified by the following specifications:

39

- [LFS] Large File Support
- [LSB] This Specification
- [SUSv2] SUSv2
- [SUSv3] ISO POSIX (2003)
- [SVID.3] SVID Issue 3
- [SVID.4] SVID Issue 4

### 13.3.1 RPC

40

#### 13.3.1.1 Interfaces for RPC

41

42

43

An LSB conforming implementation shall provide the generic functions for RPC specified in Table 13-2, with the full mandatory functionality as described in the referenced underlying specification.

44

**Table 13-2 libc - RPC Function Interfaces**

authnone_create [SVID.4]	clnt_create [SVID.4]	clnt_pcreateerror [SVID.4]	clnt_perrno [SVID.4]
clnt_perror [SVID.4]	clnt_spcreateerror [SVID.4]	clnt_sperrno [SVID.4]	clnt_s perror [SVID.4]
key_decryptsession [SVID.3]	pmap_getport [LSB]	pmap_set [LSB]	pmap_unset [LSB]
svc_getreqset [SVID.3]	svc_register [LSB]	svc_run [LSB]	svc_sendreply [LSB]
svcerr_auth [SVID.3]	svcerr_decode [SVID.3]	svcerr_noproc [SVID.3]	svcerr_noprog [SVID.3]
svcerr_progvers [SVID.3]	svcerr_systemerr [SVID.3]	svcerr_weakauth [SVID.3]	svctcp_create [LSB]
svcupd_create [LSB]	xdr_accepted_replay [SVID.3]	xdr_array [SVID.3]	xdr_bool [SVID.3]
xdr_bytes [SVID.3]	xdr_callhdr [SVID.3]	xdr_callmsg [SVID.3]	xdr_char [SVID.3]
xdr_double [SVID.3]	xdr_enum [SVID.3]	xdr_float [SVID.3]	xdr_free [SVID.3]
xdr_int [SVID.3]	xdr_long [SVID.3]	xdr_opaque [SVID.3]	xdr_opaque_auth [SVID.3]
xdr_pointer [SVID.3]	xdr_reference [SVID.3]	xdr_rejected_replay [SVID.3]	xdr_replymsg [SVID.3]
xdr_short [SVID.3]	xdr_string [SVID.3]	xdr_u_char [SVID.3]	xdr_u_int [LSB]
xdr_u_long	xdr_u_short	xdr_union	xdr_vector

45

[SVID.3]	[SVID.3]	[SVID.3]	[SVID.3]
xdr_void [SVID.3]	xdr_wrapstring [SVID.3]	xdrmem_create [SVID.3]	xdrrec_create [SVID.3]
xdrrec_eof [SVID.3]			

46

### 13.3.2 System Calls

47

#### 13.3.2.1 Interfaces for System Calls

48

49

An LSB conforming implementation shall provide the generic functions for System Calls specified in Table 13-3, with the full mandatory functionality as described in the referenced underlying specification.

50

**Table 13-3 libc - System Calls Function Interfaces**

__fxstat [LSB]	__getpgid [LSB]	__lxstat [LSB]	__xmknod [LSB]
__xstat [LSB]	access [SUSv3]	acct [LSB]	alarm [SUSv3]
brk [SUSv2]	chdir [SUSv3]	chmod [SUSv3]	chown [SUSv3]
chroot [SUSv2]	clock [SUSv3]	close [SUSv3]	closedir [SUSv3]
creat [SUSv3]	dup [SUSv3]	dup2 [SUSv3]	execl [SUSv3]
execle [SUSv3]	execlp [SUSv3]	execv [SUSv3]	execve [SUSv3]
execvp [SUSv3]	exit [SUSv3]	fchdir [SUSv3]	fchmod [SUSv3]
fchown [SUSv3]	fcntl [LSB]	fdatasync [SUSv3]	flock [LSB]
fork [SUSv3]	fstatvfs [SUSv3]	fsync [SUSv3]	ftime [SUSv3]
ftruncate [SUSv3]	getcontext [SUSv3]	getegid [SUSv3]	geteuid [SUSv3]
getgid [SUSv3]	getgroups [SUSv3]	getitimer [SUSv3]	getloadavg [LSB]
getpagesize [SUSv2]	getpgid [SUSv3]	getpgrp [SUSv3]	getpid [SUSv3]
getppid [SUSv3]	getpriority [SUSv3]	getrlimit [SUSv3]	getrusage [SUSv3]
getsid [SUSv3]	getuid [SUSv3]	getwd [SUSv3]	initgroups [LSB]
ioctl [LSB]	kill [LSB]	killpg [SUSv3]	lchown [SUSv3]
link [LSB]	lockf [SUSv3]	lseek [SUSv3]	mkdir [SUSv3]
mkfifo [SUSv3]	mlock [SUSv3]	mlockall [SUSv3]	mmap [SUSv3]
mprotect [SUSv3]	msync [SUSv3]	munlock [SUSv3]	munlockall [SUSv3]
munmap [SUSv3]	nanosleep [SUSv3]	nice [SUSv3]	open [SUSv3]

51

opendir [SUSv3]	pathconf [SUSv3]	pause [SUSv3]	pipe [SUSv3]
poll [SUSv3]	read [SUSv3]	readdir [SUSv3]	readdir_r [SUSv3]
readlink [SUSv3]	readv [SUSv3]	rename [SUSv3]	rmdir [SUSv3]
sbrk [SUSv2]	sched_get_priority_max [SUSv3]	sched_get_priority_min [SUSv3]	sched_getparam [SUSv3]
sched_getscheduler [SUSv3]	sched_rr_get_interval [SUSv3]	sched_setparam [SUSv3]	sched_setschedule [SUSv3]
sched_yield [SUSv3]	select [SUSv3]	setcontext [SUSv3]	setegid [SUSv3]
seteuid [SUSv3]	setgid [SUSv3]	setitimer [SUSv3]	setpgid [SUSv3]
setpgrp [SUSv3]	setpriority [SUSv3]	setregid [SUSv3]	setreuid [SUSv3]
setrlimit [SUSv3]	setrlimit64 [LFS]	setsid [SUSv3]	setuid [SUSv3]
sleep [SUSv3]	statvfs [SUSv3]	stime [LSB]	symlink [SUSv3]
sync [SUSv3]	sysconf [SUSv3]	time [SUSv3]	times [SUSv3]
truncate [SUSv3]	ulimit [SUSv3]	umask [SUSv3]	uname [SUSv3]
unlink [LSB]	utime [SUSv3]	utimes [SUSv3]	vfork [SUSv3]
wait [SUSv3]	wait4 [LSB]	waitpid [LSB]	write [SUSv3]
writev [SUSv3]			

### 13.3.3 Standard I/O

52

#### 13.3.3.1 Interfaces for Standard I/O

53  
54  
55

An LSB conforming implementation shall provide the generic functions for Standard I/O specified in Table 13-4, with the full mandatory functionality as described in the referenced underlying specification.

56

**Table 13-4 libc - Standard I/O Function Interfaces**

_IO_feof [LSB]	_IO_getc [LSB]	_IO_putc [LSB]	_IO_puts [LSB]
asprintf [LSB]	clearerr [SUSv3]	ctermid [SUSv3]	fclose [SUSv3]
fdopen [SUSv3]	feof [SUSv3]	ferror [SUSv3]	fflush [SUSv3]
fflush_unlocked [LSB]	fgetc [SUSv3]	fgetpos [SUSv3]	fgets [SUSv3]
fgetwc_unlocked [LSB]	fileno [SUSv3]	flockfile [SUSv3]	fopen [SUSv3]
fprintf [SUSv3]	fputc [SUSv3]	fputs [SUSv3]	fread [SUSv3]
freopen [SUSv3]	fscanf [LSB]	fseek [SUSv3]	fseeko [SUSv3]
fsetpos [SUSv3]	ftell [SUSv3]	ftello [SUSv3]	fwrite [SUSv3]
getc [SUSv3]	getc_unlocked	getchar [SUSv3]	getchar_unlocked

	[SUSv3]		[SUSv3]
getw [SUSv2]	pclose [SUSv3]	popen [SUSv3]	printf [SUSv3]
putc [SUSv3]	putc_unlocked [SUSv3]	putchar [SUSv3]	putchar_unlocked [SUSv3]
puts [SUSv3]	putw [SUSv2]	remove [SUSv3]	rewind [SUSv3]
rewinddir [SUSv3]	scanf [LSB]	seekdir [SUSv3]	setbuf [SUSv3]
setbuffer [LSB]	setvbuf [SUSv3]	snprintf [SUSv3]	sprintf [SUSv3]
sscanf [LSB]	telldir [SUSv3]	tempnam [SUSv3]	ungetc [SUSv3]
vasprintf [LSB]	vdprintf [LSB]	vfprintf [SUSv3]	vprintf [SUSv3]
vsnprintf [SUSv3]	vsprintf [SUSv3]		

57

An LSB conforming implementation shall provide the generic data interfaces for Standard I/O specified in Table 13-5, with the full mandatory functionality as described in the referenced underlying specification.

58

59

60

**Table 13-5 libc - Standard I/O Data Interfaces**

61

62

stderr [SUSv3]	stdin [SUSv3]	stdout [SUSv3]	
----------------	---------------	----------------	--

### 13.3.4 Signal Handling

63

#### 13.3.4.1 Interfaces for Signal Handling

64

65

66

An LSB conforming implementation shall provide the generic functions for Signal Handling specified in Table 13-6, with the full mandatory functionality as described in the referenced underlying specification.

67

**Table 13-6 libc - Signal Handling Function Interfaces**

__libc_current_sigrtmax [LSB]	__libc_current_sigrtmin [LSB]	__sigsetjmp [LSB]	__sysv_signal [LSB]
bsd_signal [SUSv3]	psignal [LSB]	raise [SUSv3]	sigaction [SUSv3]
sigaddset [SUSv3]	sigaltstack [SUSv3]	sigandset [LSB]	sigdelset [SUSv3]
sigemptyset [SUSv3]	sigfillset [SUSv3]	sighold [SUSv3]	sigignore [SUSv3]
siginterrupt [SUSv3]	sigisemptyset [LSB]	sigismember [SUSv3]	siglongjmp [SUSv3]
signal [SUSv3]	sigorset [LSB]	sigpause [SUSv3]	sigpending [SUSv3]
sigprocmask [SUSv3]	sigqueue [SUSv3]	sigrelse [SUSv3]	sigreturn [LSB]
sigset [SUSv3]	sigsuspend	sigtimedwait	sigwait [SUSv3]

68

	[SUSv3]	[SUSv3]	
sigwaitinfo [SUSv3]			

69

70

71

An LSB conforming implementation shall provide the generic data interfaces for Signal Handling specified in Table 13-7, with the full mandatory functionality as described in the referenced underlying specification.

72

**Table 13-7 libc - Signal Handling Data Interfaces**

73

_sys_siglist [LSB]			
--------------------	--	--	--

### 13.3.5 Localization Functions

74

#### 13.3.5.1 Interfaces for Localization Functions

75

76

77

An LSB conforming implementation shall provide the generic functions for Localization Functions specified in Table 13-8, with the full mandatory functionality as described in the referenced underlying specification.

78

**Table 13-8 libc - Localization Functions Function Interfaces**

79

bind_textdomain_codeset [LSB]	bindtextdomain [LSB]	catclose [SUSv3]	catgets [SUSv3]
catopen [SUSv3]	dcgettext [LSB]	dcngettext [LSB]	dgettext [LSB]
dngettext [LSB]	duplocale(GLIBC_2.3) [LSB]	freelocale(GLIBC_2.3) [LSB]	gettext [LSB]
iconv [SUSv3]	iconv_close [SUSv3]	iconv_open [SUSv3]	localeconv [SUSv3]
newlocale(GLIBC_2.3) [LSB]	ngettext [LSB]	nl_langinfo [SUSv3]	setlocale [SUSv3]
textdomain [LSB]	uselocale(GLIBC_2.3) [LSB]		

80

81

82

An LSB conforming implementation shall provide the generic data interfaces for Localization Functions specified in Table 13-9, with the full mandatory functionality as described in the referenced underlying specification.

83

**Table 13-9 libc - Localization Functions Data Interfaces**

84

_nl_msg_cat_cntr [LSB]			
------------------------	--	--	--

### 13.3.6 Socket Interface

85

#### 13.3.6.1 Interfaces for Socket Interface

86

87

88

An LSB conforming implementation shall provide the generic functions for Socket Interface specified in Table 13-10, with the full mandatory functionality as described in the referenced underlying specification.

89

**Table 13-10 libc - Socket Interface Function Interfaces**

<code>__h_errno_locatio n [LSB]</code>	<code>accept [SUSv3]</code>	<code>bind [SUSv3]</code>	<code>bindresvport [LSB]</code>
<code>connect [SUSv3]</code>	<code>gethostid [SUSv3]</code>	<code>gethostname [SUSv3]</code>	<code>getpeername [SUSv3]</code>
<code>getsockname [SUSv3]</code>	<code>getsockopt [LSB]</code>	<code>if_freenameindex [SUSv3]</code>	<code>if_indextoname [SUSv3]</code>
<code>if_nameindex [SUSv3]</code>	<code>if_nametoindex [SUSv3]</code>	<code>listen [SUSv3]</code>	<code>recv [SUSv3]</code>
<code>recvfrom [SUSv3]</code>	<code>recvmsg [SUSv3]</code>	<code>send [SUSv3]</code>	<code>sendmsg [SUSv3]</code>
<code>sendto [SUSv3]</code>	<code>setsockopt [LSB]</code>	<code>shutdown [SUSv3]</code>	<code>sockatmark [SUSv3]</code>
<code>socket [SUSv3]</code>	<code>socketpair [SUSv3]</code>		

90

### 13.3.7 Wide Characters

91

#### 13.3.7.1 Interfaces for Wide Characters

92

An LSB conforming implementation shall provide the generic functions for Wide Characters specified in Table 13-11, with the full mandatory functionality as described in the referenced underlying specification.

93

94

**Table 13-11 libc - Wide Characters Function Interfaces**

<code>__wcstod_internal [LSB]</code>	<code>__wcstof_internal [LSB]</code>	<code>__wcstol_internal [LSB]</code>	<code>__wcstold_interna l [LSB]</code>
<code>__wcstoul_interna l [LSB]</code>	<code>btowc [SUSv3]</code>	<code>fgetwc [SUSv3]</code>	<code>fgetws [SUSv3]</code>
<code>fputwc [SUSv3]</code>	<code>fputws [SUSv3]</code>	<code>fwide [SUSv3]</code>	<code>fwprintf [SUSv3]</code>
<code>fwscanf [LSB]</code>	<code>getwc [SUSv3]</code>	<code>getwchar [SUSv3]</code>	<code>mblen [SUSv3]</code>
<code>mbrlen [SUSv3]</code>	<code>mbrtowc [SUSv3]</code>	<code>mbsinit [SUSv3]</code>	<code>mbsnrtowcs [LSB]</code>
<code>mbsrtowcs [SUSv3]</code>	<code>mbstowcs [SUSv3]</code>	<code>mbtowc [SUSv3]</code>	<code>putwc [SUSv3]</code>
<code>putwchar [SUSv3]</code>	<code>swprintf [SUSv3]</code>	<code>swscanf [LSB]</code>	<code>towctrans [SUSv3]</code>
<code>towlower [SUSv3]</code>	<code>toupper [SUSv3]</code>	<code>ungetwc [SUSv3]</code>	<code>vfwprintf [SUSv3]</code>
<code>vfwscanf [LSB]</code>	<code>vswprintf [SUSv3]</code>	<code>vswscanf [LSB]</code>	<code>vwprintf [SUSv3]</code>
<code>vwscanf [LSB]</code>	<code>wcpncpy [LSB]</code>	<code>wcpncpy [LSB]</code>	<code>wcrtomb [SUSv3]</code>
<code>wcscasecmp [LSB]</code>	<code>wcscat [SUSv3]</code>	<code>wcschr [SUSv3]</code>	<code>wcscmp [SUSv3]</code>
<code>wcsccoll [SUSv3]</code>	<code>wcscpy [SUSv3]</code>	<code>wcscspn [SUSv3]</code>	<code>wcsdup [LSB]</code>
<code>wcsftime [SUSv3]</code>	<code>wcslen [SUSv3]</code>	<code>wcsncasecmp [LSB]</code>	<code>wcsncat [SUSv3]</code>

96

wcsncmp [SUSv3]	wcsncpy [SUSv3]	wcsnlen [LSB]	wcsnrtombs [LSB]
wcspbrk [SUSv3]	wcsrchr [SUSv3]	wcsrtombs [SUSv3]	wcsspn [SUSv3]
wcsstr [SUSv3]	wcstod [SUSv3]	wcstof [SUSv3]	wcstoiimax [SUSv3]
wcstok [SUSv3]	wcstol [SUSv3]	wcstold [SUSv3]	wcstoll [SUSv3]
wcstombs [SUSv3]	wcstoq [LSB]	wcstoul [SUSv3]	wcstoull [SUSv3]
wcstoumax [SUSv3]	wcstouq [LSB]	wcswcs [SUSv3]	wcswidth [SUSv3]
wcsxfrm [SUSv3]	wctob [SUSv3]	wctomb [SUSv3]	wctrans [SUSv3]
wctype [SUSv3]	wcwidth [SUSv3]	wmemchr [SUSv3]	wmemcmp [SUSv3]
wmemcpy [SUSv3]	wmemmove [SUSv3]	wmemset [SUSv3]	wprintf [SUSv3]
wscanf [LSB]			

### 13.3.8 String Functions

97

#### 13.3.8.1 Interfaces for String Functions

98

99

An LSB conforming implementation shall provide the generic functions for String Functions specified in Table 13-12, with the full mandatory functionality as described in the referenced underlying specification.

100

101 **Table 13-12 libc - String Functions Function Interfaces**

__memcpy [LSB]	__rawmemchr [LSB]	__stpcpy [LSB]	__strup [LSB]
__strtod_internal [LSB]	__strtodf_internal [LSB]	__strtok_r [LSB]	__strtol_internal [LSB]
__strtold_internal [LSB]	__strtoll_internal [LSB]	__strtoul_internal [LSB]	__strtoull_internal [LSB]
bcmp [SUSv3]	bcopy [SUSv3]	bzero [SUSv3]	ffs [SUSv3]
index [SUSv3]	memccpy [SUSv3]	memchr [SUSv3]	memcmp [SUSv3]
memcpy [SUSv3]	memmove [SUSv3]	memrchr [LSB]	memset [SUSv3]
rindex [SUSv3]	stpcpy [LSB]	stpcopy [LSB]	strcasecmp [SUSv3]
strcasestr [LSB]	strcat [SUSv3]	strchr [SUSv3]	strcmp [SUSv3]
strcoll [SUSv3]	strcpy [SUSv3]	strcspn [SUSv3]	strup [SUSv3]
strerror [SUSv3]	strerror_r [LSB]	strfmon [SUSv3]	strftime [SUSv3]

	strlen [SUSv3]	strncasecmp [SUSv3]	strncat [SUSv3]	strncmp [SUSv3]
102	strncpy [SUSv3]	strndup [LSB]	strnlen [LSB]	strpbrk [SUSv3]
	strptime [LSB]	strrchr [SUSv3]	strsep [LSB]	strsignal [LSB]
	strspn [SUSv3]	strstr [SUSv3]	strtof [SUSv3]	strtoimax [SUSv3]
	strtok [SUSv3]	strtok_r [SUSv3]	strtold [SUSv3]	strtoll [SUSv3]
	strtoq [LSB]	strtoull [SUSv3]	strtoumax [SUSv3]	strtouq [LSB]
	strxfrm [SUSv3]	swab [SUSv3]		

### 13.3.9 IPC Functions

#### 13.3.9.1 Interfaces for IPC Functions

An LSB conforming implementation shall provide the generic functions for IPC Functions specified in Table 13-13, with the full mandatory functionality as described in the referenced underlying specification.

107 **Table 13-13 libc - IPC Functions Function Interfaces**

ftok [SUSv3]	msgctl [SUSv3]	msgget [SUSv3]	msgrcv [SUSv3]
msgsnd [SUSv3]	semctl [SUSv3]	semget [SUSv3]	semop [SUSv3]
shmat [SUSv3]	shmctl [SUSv3]	shmdt [SUSv3]	shmget [SUSv3]

### 13.3.10 Regular Expressions

#### 13.3.10.1 Interfaces for Regular Expressions

An LSB conforming implementation shall provide the generic functions for Regular Expressions specified in Table 13-14, with the full mandatory functionality as described in the referenced underlying specification.

113 **Table 13-14 libc - Regular Expressions Function Interfaces**

regcomp [SUSv3]	regerror [SUSv3]	regexec [LSB]	regfree [SUSv3]
-----------------	------------------	---------------	-----------------

### 13.3.11 Character Type Functions

#### 13.3.11.1 Interfaces for Character Type Functions

An LSB conforming implementation shall provide the generic functions for Character Type Functions specified in Table 13-15, with the full mandatory functionality as described in the referenced underlying specification.

119 **Table 13-15 libc - Character Type Functions Function Interfaces**

__ctype_b_loc(GLIBC_2.3) [LSB]	__ctype_get_mb_cur_max [LSB]	__ctype_tolower_loc(GLIBC_2.3) [LSB]	__ctype_toupper_loc(GLIBC_2.3) [LSB]
_tolower [SUSv3]	_toupper [SUSv3]	isalnum [SUSv3]	isalpha [SUSv3]

120	isascii [SUSv3]	iscntrl [SUSv3]	isdigit [SUSv3]	isgraph [SUSv3]
	islower [SUSv3]	isprint [SUSv3]	ispunct [SUSv3]	isspace [SUSv3]
	isupper [SUSv3]	iswalnum [SUSv3]	iswalpha [SUSv3]	iswblank [SUSv3]
	iswcntrl [SUSv3]	iswctype [SUSv3]	iswdigit [SUSv3]	iswgraph [SUSv3]
	iswlower [SUSv3]	iswprint [SUSv3]	iswpunct [SUSv3]	iswspace [SUSv3]
	iswupper [SUSv3]	iswxdigit [SUSv3]	isxdigit [SUSv3]	toascii [SUSv3]
	tolower [SUSv3]	toupper [SUSv3]		

### 13.3.12 Time Manipulation

#### 13.3.12.1 Interfaces for Time Manipulation

An LSB conforming implementation shall provide the generic functions for Time Manipulation specified in Table 13-16, with the full mandatory functionality as described in the referenced underlying specification.

125 **Table 13-16 libc - Time Manipulation Function Interfaces**

126	adjtime [LSB]	asctime [SUSv3]	asctime_r [SUSv3]	ctime [SUSv3]
	ctime_r [SUSv3]	difftime [SUSv3]	gmtime [SUSv3]	gmtime_r [SUSv3]
	localtime [SUSv3]	localtime_r [SUSv3]	mktime [SUSv3]	tzset [SUSv3]
	ualarm [SUSv3]			

127 An LSB conforming implementation shall provide the generic data interfaces for  
 128 Time Manipulation specified in Table 13-17, with the full mandatory functionality as  
 129 described in the referenced underlying specification.

130 **Table 13-17 libc - Time Manipulation Data Interfaces**

131	__daylight [LSB]	__timezone [LSB]	__tzname [LSB]	daylight [SUSv3]
	timezone [SUSv3]	tzname [SUSv3]		

### 13.3.13 Terminal Interface Functions

#### 13.3.13.1 Interfaces for Terminal Interface Functions

An LSB conforming implementation shall provide the generic functions for Terminal Interface Functions specified in Table 13-18, with the full mandatory functionality as described in the referenced underlying specification.

136 **Table 13-18 libc - Terminal Interface Functions Function Interfaces**

cfgetispeed [SUSv3]	cfgetospeed [SUSv3]	cfmakeraw [LSB]	cfsetispeed [SUSv3]
cfsetospeed [SUSv3]	cfsetspeed [LSB]	tcdrain [SUSv3]	tcflow [SUSv3]
tcflush [SUSv3]	tcgetattr [SUSv3]	tcgetpgrp [SUSv3]	tcgetsid [SUSv3]

137

tcsetattr [SUSv3]	tcsetattr [SUSv3]	tcsetpgroup [SUSv3]	
-------------------	-------------------	---------------------	--

138

### 13.3.14 System Database Interface

139

140

141

#### 13.3.14.1 Interfaces for System Database Interface

An LSB conforming implementation shall provide the generic functions for System Database Interface specified in Table 13-19, with the full mandatory functionality as described in the referenced underlying specification.

142

**Table 13-19 libc - System Database Interface Function Interfaces**

endgrent [SUSv3]	endprotoent [SUSv3]	endpwent [SUSv3]	endservent [SUSv3]
endutent [SUSv2]	endutxent [SUSv3]	getgrent [SUSv3]	getgrgid [SUSv3]
getrgid_r [SUSv3]	getgrnam [SUSv3]	getgrnam_r [SUSv3]	getgrouplist [LSB]
gethostbyaddr [SUSv3]	gethostbyname [SUSv3]	getprotobynamename [SUSv3]	getprotobyname [SUSv3]
getprotoent [SUSv3]	getpwent [SUSv3]	getpwnam [SUSv3]	getpwnam_r [SUSv3]
getpwuid [SUSv3]	getpwuid_r [SUSv3]	getservbyname [SUSv3]	getservbyport [SUSv3]
getservent [SUSv3]	getutent [LSB]	getutent_r [LSB]	getutxent [SUSv3]
getutxid [SUSv3]	getutxline [SUSv3]	pututxline [SUSv3]	setgrent [SUSv3]
setgroups [LSB]	setprotoent [SUSv3]	setpwent [SUSv3]	setservent [SUSv3]
setutent [LSB]	setutxent [SUSv3]	utmpname [LSB]	

143

### 13.3.15 Language Support

144

145

146

147

#### 13.3.15.1 Interfaces for Language Support

An LSB conforming implementation shall provide the generic functions for Language Support specified in Table 13-20, with the full mandatory functionality as described in the referenced underlying specification.

148

**Table 13-20 libc - Language Support Function Interfaces**

__libc_start_main [LSB]	__register_atfork( GLIBC_2.3.2) [LSB]		
-------------------------	---------------------------------------	--	--

149

### 13.3.16 Large File Support

150

#### 13.3.16.1 Interfaces for Large File Support

151

152

153

An LSB conforming implementation shall provide the generic functions for Large File Support specified in Table 13-21, with the full mandatory functionality as described in the referenced underlying specification.

154

**Table 13-21 libc - Large File Support Function Interfaces**

__fxstat64 [LSB]	__lxstat64 [LSB]	__xstat64 [LSB]	creat64 [LFS]
fgetpos64 [LFS]	fopen64 [LFS]	freopen64 [LFS]	fseeko64 [LFS]
fsetpos64 [LFS]	fstatvfs64 [LFS]	ftello64 [LFS]	ftruncate64 [LFS]
ftw64 [LFS]	getrlimit64 [LFS]	lockf64 [LFS]	mkstemp64 [LFS]
mmap64 [LFS]	nftw64 [LFS]	readdir64 [LFS]	statvfs64 [LFS]
tmpfile64 [LFS]	truncate64 [LFS]		

155

### 13.3.17 Standard Library

156

#### 13.3.17.1 Interfaces for Standard Library

157

158

159

An LSB conforming implementation shall provide the generic functions for Standard Library specified in Table 13-22, with the full mandatory functionality as described in the referenced underlying specification.

160

**Table 13-22 libc - Standard Library Function Interfaces**

_Exit [SUSv3]	__assert_fail [LSB]	__cxa_atexit [LSB]	__errno_location [LSB]
__fpending [LSB]	__getpagesize [LSB]	__isinf [LSB]	__isinff [LSB]
__isinf [LSB]	__isnan [LSB]	__isnanf [LSB]	__isnanl [LSB]
__sysconf [LSB]	_exit [SUSv3]	_longjmp [SUSv3]	_setjmp [SUSv3]
a64l [SUSv3]	abort [SUSv3]	abs [SUSv3]	atof [SUSv3]
atoi [SUSv3]	atol [SUSv3]	atoll [SUSv3]	basename [SUSv3]
bsearch [SUSv3]	calloc [SUSv3]	closelog [SUSv3]	confstr [SUSv3]
cuserid [SUSv2]	daemon [LSB]	dirname [SUSv3]	div [SUSv3]
drand48 [SUSv3]	ecvt [SUSv3]	erand48 [SUSv3]	err [LSB]
error [LSB]	errx [LSB]	fcvt [SUSv3]	fmtmsg [SUSv3]
fnmatch [SUSv3]	fpathconf [SUSv3]	free [SUSv3]	freeaddrinfo [SUSv3]
ftrylockfile [SUSv3]	ftw [SUSv3]	funlockfile [SUSv3]	gai_strerror [SUSv3]
gcvt [SUSv3]	getaddrinfo [SUSv3]	getcwd [SUSv3]	getdate [SUSv3]

getenv [SUSv3]	getlogin [SUSv3]	getlogin_r [SUSv3]	getnameinfo [SUSv3]
getopt [LSB]	getopt_long [LSB]	getopt_long_only [LSB]	getsubopt [SUSv3]
gettimeofday [SUSv3]	glob [SUSv3]	glob64 [LSB]	globfree [SUSv3]
globfree64 [LSB]	grantpt [SUSv3]	hcreate [SUSv3]	hdestroy [SUSv3]
hsearch [SUSv3]	htonl [SUSv3]	htons [SUSv3]	imaxabs [SUSv3]
imaxdiv [SUSv3]	inet_addr [SUSv3]	inet_ntoa [SUSv3]	inet_ntop [SUSv3]
inet_pton [SUSv3]	initstate [SUSv3]	insque [SUSv3]	isatty [SUSv3]
isblank [SUSv3]	jrand48 [SUSv3]	l64a [SUSv3]	labs [SUSv3]
lcong48 [SUSv3]	ldiv [SUSv3]	lfind [SUSv3]	llabs [SUSv3]
lldiv [SUSv3]	longjmp [SUSv3]	lrand48 [SUSv3]	lsearch [SUSv3]
makecontext [SUSv3]	malloc [SUSv3]	memmem [LSB]	mkstemp [SUSv3]
mktemp [SUSv3]	mrand48 [SUSv3]	nftw [SUSv3]	nrand48 [SUSv3]
ntohl [SUSv3]	ntohs [SUSv3]	openlog [SUSv3]	perror [SUSv3]
posix_memalign [SUSv3]	posix_openpt [SUSv3]	ptsname [SUSv3]	putenv [SUSv3]
qsort [SUSv3]	rand [SUSv3]	rand_r [SUSv3]	random [SUSv3]
realloc [SUSv3]	realpath [SUSv3]	remque [SUSv3]	seed48 [SUSv3]
setenv [SUSv3]	sethostname [LSB]	setlogmask [SUSv3]	setstate [SUSv3]
srand [SUSv3]	srand48 [SUSv3]	srandom [SUSv3]	strtod [SUSv3]
strtol [SUSv3]	strtoul [SUSv3]	swapcontext [SUSv3]	syslog [SUSv3]
system [LSB]	tdelete [SUSv3]	tfind [SUSv3]	tmpfile [SUSv3]
tmpnam [SUSv3]	tsearch [SUSv3]	ttynname [SUSv3]	ttynname_r [SUSv3]
twalk [SUSv3]	unlockpt [SUSv3]	unsetenv [SUSv3]	usleep [SUSv3]
verrx [LSB]	vfscanf [LSB]	vscanf [LSB]	vsscanf [LSB]
vsyslog [LSB]	warn [LSB]	warnx [LSB]	wordexp [SUSv3]
wordfree [SUSv3]			

161

An LSB conforming implementation shall provide the generic data interfaces for Standard Library specified in Table 13-23, with the full mandatory functionality as described in the referenced underlying specification.

162  
163  
164

165

**Table 13-23 libc - Standard Library Data Interfaces**

166

_environ [LSB]	_environ [LSB]	_sys_errlist [LSB]	environ [SUSv3]
getdate_err [SUSv3]	optarg [SUSv3]	opterr [SUSv3]	optind [SUSv3]
optopt [SUSv3]			

## 13.4 Data Definitions for libc

167

This section defines global identifiers and their values that are associated with interfaces contained in libc. These definitions are organized into groups that correspond to system headers. This convention is used as a convenience for the reader, and does not imply the existence of these headers, or their content. Where an interface is defined as requiring a particular system header file all of the data definitions for that system header file presented here shall be in effect.

168

169

170

171

172

173

174

175

176

This section gives data definitions to promote binary application portability, not to repeat source interface definitions available elsewhere. System providers and application developers should use this ABI to supplement - not to replace - source interface definition specifications.

177

178

179

180

This specification uses the ISO C (1999) C Language as the reference programming language, and data definitions are specified in ISO C format. The C language is used here as a convenient notation. Using a C language description of these data objects does not preclude their use by other programming languages.

### 13.4.1 arpa/inet.h

181

182

183

184

185

186

187

188

189

```
extern uint32_t htonl(uint32_t);
extern uint16_t htons(uint16_t);
extern in_addr_t inet_addr(const char *);
extern char *inet_ntoa(struct in_addr);
extern const char *inet_ntop(int, const void *, char *, socklen_t);
extern int inet_pton(int, const char *, void *);
extern uint32_t ntohl(uint32_t);
extern uint16_t ntohs(uint16_t);
```

### 13.4.2 assert.h

190

191

192

193

The assert.h header shall define the assert() macro. It refers to the macro NDEBUG, which is not defined in this header. If NDEBUG is defined before the inclusion of this header, the assert() macro shall be defined as described below, otherwise the macro shall behave as described in assert() in ISO/IEC 9945 POSIX.

194

195

196

```
extern void __assert_fail(const char *, const char *, unsigned int,
                         const char *);
```

### 13.4.3 ctype.h

197

198

199

200

201

202

203

```
enum {
    _ISupper, _ISlower, _ISalpha, _ISdigit, _ISxdigit, _ISspace,
    _ISprint,
    _ISgraph, _ISblank, _IScntrl, _ISPunct, _ISalnum
};
extern int _tolower(int);
```

```

204     extern int _toupper(int);
205     extern int isalnum(int);
206     extern int isalpha(int);
207     extern int isascii(int);
208     extern int iscntrl(int);
209     extern int isdigit(int);
210     extern int isgraph(int);
211     extern int islower(int);
212     extern int isprint(int);
213     extern int ispunct(int);
214     extern int isspace(int);
215     extern int isupper(int);
216     extern int isxdigit(int);
217     extern int toascii(int);
218     extern int tolower(int);
219     extern int toupper(int);
220     extern int isblank(int);
221     extern const unsigned short **__ctype_b_loc(void);
222     extern const int32_t **__ctype_toupper_loc(void);
223     extern const int32_t **__ctype_tolower_loc(void);

```

#### 13.4.4 dirent.h

```

224     typedef struct __dirstream DIR;
225
226     struct dirent {
227         long int d_ino;
228         off_t d_off;
229         unsigned short d_reclen;
230         unsigned char d_type;
231         char d_name[256];
232     };
233     struct dirent64 {
234         uint64_t d_ino;
235         int64_t d_off;
236         unsigned short d_reclen;
237         unsigned char d_type;
238         char d_name[256];
239     };
240     extern void rewinddir(DIR *);
241     extern void seekdir(DIR *, long int);
242     extern long int telldir(DIR *);
243     extern int closedir(DIR *);
244     extern DIR *opendir(const char *);
245     extern struct dirent *readdir(DIR *);
246     extern struct dirent64 *readdir64(DIR *);
247     extern int readdir_r(DIR *, struct dirent *, struct dirent **);
248

```

#### 13.4.5 err.h

```

249     extern void err(int, const char *, ...);
250     extern void errx(int, const char *, ...);
251     extern void warn(const char *, ...);
252     extern void warnx(const char *, ...);
253     extern void error(int, int, const char *, ...);
254

```

#### 13.4.6 errno.h

ISO POSIX (2003) requires that each error value shall be unique, with permission for EAGAIN and EWOULDBLOCK possibly having the same value. This specification also requires that ENOTSUP and EOPNOTSUPP have the same value.

258           **Note:** A defect report against ISO POSIX (2003) has been filed to request that  
 259           specification also permit these two symbols to have the same value.

```

260
261     #define errno    (*__errno_location())
262
263     #define EPERM     1
264     #define ECHILD    10
265     #define ENETDOWN   100
266     #define ENETUNREACH 101
267     #define ENETRESET   102
268     #define ECONNABORTED 103
269     #define ECONNRESET   104
270     #define ENOBUFS   105
271     #define EISCONN   106
272     #define ENOTCONN  107
273     #define ESHUTDOWN 108
274     #define ETOOMANYREFS 109
275     #define EAGAIN    11
276     #define ETIMEDOUT 110
277     #define ECONNREFUSED 111
278     #define EHOSTDOWN  112
279     #define EHOSTUNREACH 113
280     #define EALREADY   114
281     #define EINPROGRESS 115
282     #define ESTALE    116
283     #define EUCLEAN   117
284     #define ENOTNAM   118
285     #define ENAVAIL   119
286     #define ENOMEM    12
287     #define EISNAM    120
288     #define EREMOTEIO 121
289     #define EDQUOT   122
290     #define ENOMEDIUM 123
291     #define EMEDIUMTYPE 124
292     #define ECANCELED 125
293     #define EACCES   13
294     #define EFAULT   14
295     #define ENOTBLK   15
296     #define EBUSY    16
297     #define EEXIST   17
298     #define EXDEV    18
299     #define ENODEV   19
300     #define ENOENT   2
301     #define ENOTDIR  20
302     #define EISDIR   21
303     #define EINVAL   22
304     #define ENFILE   23
305     #define EMFILE   24
306     #define ENOTTY   25
307     #define ETXTBSY  26
308     #define EFBIG    27
309     #define ENOSPC   28
310     #define ESPIPE   29
311     #define ESRCH    3
312     #define EROFS   30
313     #define EMLINK   31
314     #define EPIPE    32
315     #define EDOM    33
316     #define ERANGE   34
317     #define EDEADLK  35
318     #define ENAMETOOLONG 36
319     #define ENOLCK   37
320     #define ENOSYS   38

```

```

321 #define ENOTEMPTY      39
322 #define EINTR       4
323 #define ELOOP       40
324 #define ENOMSG      42
325 #define EIDRM       43
326 #define ECHRNG      44
327 #define EL2NSYNC     45
328 #define EL3HLT      46
329 #define EL3RST      47
330 #define ELNRNG      48
331 #define EUNATCH     49
332 #define EIO        5
333 #define ENOANO      55
334 #define EBADRQC     56
335 #define EBADSLT     57
336 #define EBFONT      59
337 #define ENXIO       6
338 #define ENOSTR      60
339 #define ENODATA     61
340 #define ETIME       62
341 #define ENOSR       63
342 #define ENONET      64
343 #define ENOPKG      65
344 #define EREMOTE     66
345 #define ENOLINK     67
346 #define EADV        68
347 #define ESRMNT     69
348 #define E2BIG        7
349 #define ECOMM       70
350 #define EPROTO      71
351 #define EMULTIHOP    72
352 #define EDOTDOT     73
353 #define EBADMSG     74
354 #define EOVERRLOW   75
355 #define ENOTUNIQ    76
356 #define EBADFD      77
357 #define EREMCHG     78
358 #define ELIBACC     79
359 #define ENOEXEC      8
360 #define ELIBBAD     80
361 #define ELIBSCN     81
362 #define ELIBMAX     82
363 #define ELIBEXEC    83
364 #define EILSEQ       84
365 #define ERESTART    85
366 #define ESTRPIPE    86
367 #define EUSERS      87
368 #define ENOTSOCK    88
369 #define EDESTADDRREQ 89
370 #define EBADF       9
371 #define EMSGSIZE    90
372 #define EPROTOTYPE  91
373 #define ENOPROTOOPT 92
374 #define EPROTONOSUPPORT 93
375 #define ESOCKTNOSUPPORT 94
376 #define EOPNOTSUPP  95
377 #define EPFNOSUPPORT 96
378 #define EAFNOSUPPORT 97
379 #define EADDRINUSE   98
380 #define EADDRNOTAVAIL 99
381 #define EWOULDBLOCK  EAGAIN
382 #define ENOTSUP     EOPNOTSUPP
383
384 extern int *__errno_location(void);

```

### 13.4.7 fcntl.h

```

385     #define O_RDONLY      00
386     #define O_ACCMODE      0003
387     #define O_WRONLY       01
388     #define O_CREAT        0100
389     #define O_TRUNC        01000
390     #define O_SYNC         010000
391     #define O_RDWR         02
392     #define O_EXCL        0200
393     #define O_APPEND       020000
394     #define O_ASYNC        0200000
395     #define O_NOCTTY       0400
396     #define O_NDELAY        04000
397     #define O_NONBLOCK      04000
398     #define FD_CLOEXEC     1
399
400
401     struct flock {
402         short l_type;
403         short l_whence;
404         off_t l_start;
405         off_t l_len;
406         pid_t l_pid;
407     };
408     struct flock64 {
409         short l_type;
410         short l_whence;
411         loff_t l_start;
412         loff_t l_len;
413         pid_t l_pid;
414     };
415
416     #define F_DUPFD    0
417     #define F_RDLCK    0
418     #define F_GETFD    1
419     #define F_WRLCK    1
420     #define F_SETFD    2
421     #define F_UNLCK    2
422     #define F_GETFL    3
423     #define F_SETFL    4
424     #define F_GETLK    5
425     #define F_SETLK    6
426     #define F_SETLKW   7
427     #define F_SETOWN   8
428     #define F_GETOWN   9
429
430     extern int lockf64(int, int, off64_t);
431     extern int fcntl(int, int, ...);

```

### 13.4.8 fmtmsg.h

```

432
433     #define MM_HARD    1
434     #define MM_NRECOV   128
435     #define MM_UTIL    16
436     #define MM_SOFT    2
437     #define MM_OPSYS   32
438     #define MM_FIRM    4
439     #define MM_RECOVER  64
440     #define MM_APPL    8
441
442     #define MM_NOSEV   0
443     #define MM_HALT    1

```

```

444 #define MM_ERROR          2
445
446 #define MM_NULLLBL        ((char *) 0)
447
448 extern int fmtmsg(long int, const char *, int, const char *, const char
449 *,
450           const char *);

```

### 13.4.9 fnmatch.h

```

451
452 #define FNM_PATHNAME      (1<<0)
453 #define FNM_NOESCAPE      (1<<1)
454 #define FNM_PERIOD         (1<<2)
455 #define FNM_NOMATCH       1
456
457 extern int fnmatch(const char *, const char *, int);

```

### 13.4.10 ftw.h

```

458
459 #define FTW_D    FTW_D
460 #define FTW_DNR  FTW_DNR
461 #define FTW_DP   FTW_DP
462 #define FTW_F    FTW_F
463 #define FTW_NS   FTW_NS
464 #define FTW_SL   FTW_SL
465 #define FTW_SLN  FTW_SLN
466
467 enum {
468     FTW_F, FTW_D, FTW_DNR, FTW_NS, FTW_SL, FTW_DP, FTW_SLN
469 };
470
471 enum {
472     FTW_PHYS, FTW_MOUNT, FTW_CHDIR, FTW_DEPTH
473 };
474
475 struct FTW {
476     int base;
477     int level;
478 };
479
480 typedef int (*__ftw_func_t) (char *__filename, struct stat * __status,
481                             int __flag);
482 typedef int (*__ftw64_func_t) (char *__filename, struct stat64 * __status,
483                               int __flag);
484 typedef int (*__nftw_func_t) (char *__filename, struct stat * __status,
485                             int __flag, struct FTW * __info);
486 typedef int (*__nftw64_func_t) (char *__filename, struct stat64 * __status,
487                               int __flag, struct FTW * __info);
488
489 extern int ftw(const char *, __ftw_func_t, int);
490 extern int ftw64(const char *, __ftw64_func_t, int);
491 extern int nftw(const char *, __nftw_func_t, int, int);
492 extern int nftw64(const char *, __nftw64_func_t, int, int);
493

```

### 13.4.11 getopt.h

```

494
495 #define no_argument      0
496 #define required_argument 1
497 #define optional_argument 2

```

```

498     struct option {
499         char *name;
500         int has_arg;
501         int *flag;
502         int val;
503     };
504     extern int getopt_long(int, char *const, const char *,
505                           const struct option *, int *);
506     extern int getopt_long_only(int, char *const, const char *,
507                               const struct option *, int *);
508

```

### 13.4.12 glob.h

```

509 #define GLOB_ERR          (1<<0)
510 #define GLOB_MARK         (1<<1)
511 #define GLOB_BRACE        (1<<10)
512 #define GLOB_NOMAGIC      (1<<11)
513 #define GLOB_TILDE        (1<<12)
514 #define GLOB_ONLYDIR      (1<<13)
515 #define GLOB_TILDE_CHECK  (1<<14)
516 #define GLOB_NOSORT        (1<<2)
517 #define GLOB_DOOFFS       (1<<3)
518 #define GLOB_NOCHECK      (1<<4)
519 #define GLOB_APPEND        (1<<5)
520 #define GLOB_NOESCAPE     (1<<6)
521 #define GLOB_PERIOD        (1<<7)
522 #define GLOB_MAGCHAR      (1<<8)
523 #define GLOB_ALTDIRFUNC   (1<<9)
524
525 #define GLOB_NOSPACE       1
526 #define GLOB_ABORTED       2
527 #define GLOB_NOMATCH      3
528 #define GLOB_NOSYS         4
529
530
531     typedef struct {
532         size_t gl_pathc;
533         char **gl_pathv;
534         size_t gl_offs;
535         int gl_flags;
536         void (*gl_closedir) (void *);
537         struct dirent *(*gl_readdir) (void *);
538         void *(*gl_opendir) (const char *);
539         int (*gl_lstat) (const char *, struct stat *);
540         int (*gl_stat) (const char *, struct stat *);
541     } glob_t;
542
543     typedef struct {
544         size_t gl_pathc;
545         char **gl_pathv;
546         size_t gl_offs;
547         int gl_flags;
548         void (*gl_closedir) (void *);
549         struct dirent64 *(*gl_readdir64) (void *);
550         void *(*gl_opendir) (const char *);
551         int (*gl_lstat) (const char *, struct stat *);
552         int (*gl_stat) (const char *, struct stat *);
553     } glob64_t;
554
555     extern int glob(const char *, int,
556                    int (*__errfunc) (const char *p1, int p2)
557                    , glob_t *);
558     extern int glob64(const char *, int,
559                      int (*__errfunc) (const char *p1, int p2)

```

```

559 , glob64_t *);
560 extern void globfree(glob_t *);
561 extern void globfree64(glob64_t *);

```

### 13.4.13 grp.h

```

562
563     struct group {
564         char *gr_name;
565         char *gr_passwd;
566         gid_t gr_gid;
567         char **gr_mem;
568     };
569
570     extern void endgrent(void);
571     extern struct group *getgrent(void);
572     extern struct group *getgrgid(gid_t);
573     extern struct group *getgrnam(char *);
574     extern int initgroups(const char *, gid_t);
575     extern void setgrent(void);
576     extern int setgroups(size_t, const gid_t *);
577     extern int getgrgid_r(gid_t, struct group *, char *, size_t,
578                           struct group **);
579     extern int getgrnam_r(const char *, struct group *, char *, size_t,
580                           struct group **);
580     extern int getgrouplist(const char *, gid_t, gid_t *, int *);

```

### 13.4.14 iconv.h

```

582
583     typedef void *iconv_t;
584     extern size_t iconv(iconv_t, char **, size_t *, char **, size_t *);
585     extern int iconv_close(iconv_t);
586     extern iconv_t iconv_open(char *, char *);

```

### 13.4.15 inttypes.h

```

587
588     typedef lldiv_t imaxdiv_t;
589     typedef unsigned char uint8_t;
590     typedef unsigned short uint16_t;
591     typedef unsigned int uint32_t;
592
593     extern intmax_t strtoimax(const char *, char **, int);
594     extern uintmax_t strtoumax(const char *, char **, int);
595     extern intmax_t wcstoimax(const wchar_t *, wchar_t **, int);
596     extern uintmax_t wcstoumax(const wchar_t *, wchar_t **, int);
597     extern intmax_t imaxabs(intmax_t);
598     extern imaxdiv_t imaxdiv(intmax_t, intmax_t);

```

### 13.4.16 langinfo.h

```

599
600     #define ABDAY_1 0x20000
601     #define ABDAY_2 0x20001
602     #define ABDAY_3 0x20002
603     #define ABDAY_4 0x20003
604     #define ABDAY_5 0x20004
605     #define ABDAY_6 0x20005
606     #define ABDAY_7 0x20006
607
608     #define DAY_1    0x20007
609     #define DAY_2    0x20008

```

```

610      #define DAY_3    0x20009
611      #define DAY_4    0x2000A
612      #define DAY_5    0x2000B
613      #define DAY_6    0x2000C
614      #define DAY_7    0x2000D
615
616      #define ABMON_1   0x2000E
617      #define ABMON_2   0x2000F
618      #define ABMON_3   0x20010
619      #define ABMON_4   0x20011
620      #define ABMON_5   0x20012
621      #define ABMON_6   0x20013
622      #define ABMON_7   0x20014
623      #define ABMON_8   0x20015
624      #define ABMON_9   0x20016
625      #define ABMON_10  0x20017
626      #define ABMON_11  0x20018
627      #define ABMON_12  0x20019
628
629      #define MON_1     0x2001A
630      #define MON_2     0x2001B
631      #define MON_3     0x2001C
632      #define MON_4     0x2001D
633      #define MON_5     0x2001E
634      #define MON_6     0x2001F
635      #define MON_7     0x20020
636      #define MON_8     0x20021
637      #define MON_9     0x20022
638      #define MON_10    0x20023
639      #define MON_11    0x20024
640      #define MON_12    0x20025
641
642      #define AM_STR    0x20026
643      #define PM_STR    0x20027
644
645      #define D_T_FMT   0x20028
646      #define D_FMT     0x20029
647      #define T_FMT     0x2002A
648      #define T_FMT_AMPM 0x2002B
649
650      #define ERA       0x2002C
651      #define ERA_D_FMT 0x2002E
652      #define ALT_DIGITS 0x2002F
653      #define ERA_D_T_FMT 0x20030
654      #define ERA_T_FMT  0x20031
655
656      #define CODESET 14
657
658      #define CRNCYSTR   0x4000F
659
660      #define RADIXCHAR  0x10000
661      #define THOUSEP   0x10001
662      #define YESEXPR    0x50000
663      #define NOEXPR     0x50001
664      #define YESSTR    0x50002
665      #define NOSTR     0x50003
666
667      extern char *nl_langinfo(nl_item);

```

### 13.4.17 libgen.h

```

668
669      extern char *basename(const char *);
670      extern char *dirname(char *);

```

### 13.4.18 libintl.h

```

671     extern char *bindtextdomain(const char *, const char *);
672     extern char *dcgettext(const char *, const char *, int);
673     extern char *dgettext(const char *, const char *);
674     extern char *gettext(const char *);
675     extern char *textdomain(const char *);
676     extern char *bind_textdomain_codeset(const char *, const char *);
677     extern char *dcngettext(const char *, const char *, const char *,
678                             unsigned long int, int);
679     extern char *dngettext(const char *, const char *, const char *,
680                           unsigned long int);
681     extern char *ngettext(const char *, const char *, unsigned long int);
682

```

### 13.4.19 limits.h

```

683     #define LLONG_MIN          (-LLONG_MAX-1LL)
684     #define ULLONG_MAX         18446744073709551615ULL
685     #define OPEN_MAX           256
686     #define PATH_MAX           4096
687     #define LLONG_MAX          9223372036854775807LL
688     #define SSIZE_MAX          LONG_MAX
689
690     #define MB_LEN_MAX          16
691
692     #define SCHAR_MIN           (-128)
693     #define SCHAR_MAX           127
694     #define UCHAR_MAX           255
695     #define CHAR_BIT             8
696
697     #define SHRT_MIN            (-32768)
698     #define SHRT_MAX             32767
699     #define USHRT_MAX            65535
700
701     #define INT_MIN             (-INT_MAX-1)
702     #define INT_MAX              2147483647
703     #define __INT_MAX__          2147483647
704     #define UINT_MAX              4294967295U
705
706     #define LONG_MIN             (-LONG_MAX-1L)
707
708     #define PTHREAD_KEYS_MAX      1024
709     #define PTHREAD_THREADS_MAX    16384
710     #define PTHREAD_DESTRUCTOR_ITERATIONS 4
711

```

### 13.4.20 locale.h

```

712     struct lconv {
713         char *decimal_point;
714         char *thousands_sep;
715         char *grouping;
716         char *int_curr_symbol;
717         char *currency_symbol;
718         char *mon_decimal_point;
719         char *mon_thousands_sep;
720         char *mon_grouping;
721         char *positive_sign;
722         char *negative_sign;
723         char int_frac_digits;
724         char frac_digits;
725

```

```

726         char p_cs_precedes;
727         char p_sep_by_space;
728         char n_cs_precedes;
729         char n_sep_by_space;
730         char p_sign_posn;
731         char n_sign_posn;
732         char int_p_cs_precedes;
733         char int_p_sep_by_space;
734         char int_n_cs_precedes;
735         char int_n_sep_by_space;
736         char int_p_sign_posn;
737         char int_n_sign_posn;
738     };
739
740 #define LC_GLOBAL_LOCALE      ((locale_t) -1L)
741 #define LC_CTYPE              0
742 #define LC_NUMERIC             1
743 #define LC_TELEPHONE           10
744 #define LC_MEASUREMENT         11
745 #define LC_IDENTIFICATION      12
746 #define LC_TIME                2
747 #define LC_COLLATE              3
748 #define LC_MONETARY             4
749 #define LC_MESSAGES              5
750 #define LC_ALL                  6
751 #define LC_PAPER                7
752 #define LC_NAME                 8
753 #define LC_ADDRESS               9
754
755 typedef struct __locale_struct {
756     struct locale_data *__locales[13];
757     const unsigned short *__ctype_b;
758     const int *__ctype_tolower;
759     const int *__ctype_toupper;
760     const char *__names[13];
761 } *__locale_t;
762
763 typedef struct __locale_struct *locale_t;
764
765 #define LC_ADDRESS_MASK (1 << LC_ADDRESS)
766 #define LC_COLLATE_MASK (1 << LC_COLLATE)
767 #define LC_IDENTIFICATION_MASK (1 << LC_IDENTIFICATION)
768 #define LC_MEASUREMENT_MASK (1 << LC_MEASUREMENT)
769 #define LC_MESSAGES_MASK (1 << LC_MESSAGES)
770 #define LC_MONETARY_MASK (1 << LC_MONETARY)
771 #define LC_NAME_MASK (1 << LC_NAME)
772 #define LC_NUMERIC_MASK (1 << LC_NUMERIC)
773 #define LC_PAPER_MASK (1 << LC_PAPER)
774 #define LC_TELEPHONE_MASK (1 << LC_TELEPHONE)
775 #define LC_TIME_MASK (1 << LC_TIME)
776 #define LC_CTYPE_MASK (1<<LC_CTYPE)
777 #define LC_ALL_MASK \
778     (LC_CTYPE_MASK | LC_NUMERIC_MASK | LC_TIME_MASK | \
779     LC_COLLATE_MASK | LC_MONETARY_MASK | \
780     LC_MESSAGES_MASK | LC_PAPER_MASK | LC_NAME_MASK | \
781     LC_ADDRESS_MASK | LC_TELEPHONE_MASK | \
782     LC_MEASUREMENT_MASK | LC_IDENTIFICATION_MASK)
783
784 extern struct lconv *localeconv(void);
785 extern char *setlocale(int, const char *);
786 extern locale_t uselocale(locale_t);
787 extern void freelocale(locale_t);
788 extern locale_t duplocale(locale_t);
789 extern locale_t newlocale(int, const char *, locale_t);

```

### 13.4.21 monetary.h

```
790
791     extern ssize_t strfmon(char *, size_t, const char *, ...);
```

### 13.4.22 net/if.h

```
792
793 #define IF_NAMESIZE      16
794
795 #define IFF_UP    0x01
796 #define IFF_BROADCAST 0x02
797 #define IFF_DEBUG   0x04
798 #define IFF_LOOPBACK 0x08
799 #define IFF_POINTOPOINT 0x10
800 #define IFF_PROMISC  0x100
801 #define IFF_MULTICAST 0x1000
802 #define IFF_NOTRAILERS 0x20
803 #define IFF_RUNNING   0x40
804 #define IFF_NOARP    0x80
805
806     struct if_nameindex {
807         unsigned int if_index;
808         char *if_name;
809     };
810
811     struct ifaddr {
812         struct sockaddr ifa_addr;
813         union {
814             struct sockaddr ifu_broadaddr;
815             struct sockaddr ifu_dstaddr;
816         } ifa_ifu;
817         void *ifa_ifp;
818         void *ifa_next;
819     };
820
821 #define IFNAMSIZ        IF_NAMESIZE
822
823     struct ifreq {
824         union {
825             char ifrn_name[IFNAMSIZ];
826         } ifr_ifrn;
827         union {
828             struct sockaddr ifru_addr;
829             struct sockaddr ifru_dstaddr;
830             struct sockaddr ifru_broadaddr;
831             struct sockaddr ifru_netmask;
832             struct sockaddr ifru_hwaddr;
833             short ifru_flags;
834             int ifru_ivalue;
835             int ifru_mtu;
836             char ifru_slave[IFNAMSIZ];
837             char ifru_newname[IFNAMSIZ];
838             caddr_t ifru_data;
839             struct ifmap ifru_map;
840         } ifr_ifru;
841     };
842
843     struct ifconf {
844         int ifc_len;
845         union {
846             caddr_t ifcu_buf;
847             struct ifreq *ifcu_req;
848         } ifc_ifcu;
```

```

849      };
850      extern void if_freenameindex(struct if_nameindex *);
851      extern char *if_indextoname(unsigned int, char *);
852      extern struct if_nameindex *if_nameindex(void);
853      extern unsigned int if_nametoindex(const char *);

```

### 13.4.23 netdb.h

```

854
855      #define NETDB_INTERNAL -1
856      #define NETDB_SUCCESS 0
857      #define HOST_NOT_FOUND 1
858      #define IPPORT_RESERVED 1024
859      #define NI_MAXHOST 1025
860      #define TRY AGAIN 2
861      #define NO_RECOVERY 3
862      #define NI_MAXSERV 32
863      #define NO_DATA 4
864      #define h_addr h_addr_list[0]
865      #define NO_ADDRESS NO_DATA
866
867      struct servent {
868          char *s_name;
869          char **s_aliases;
870          int s_port;
871          char *s_proto;
872      };
873      struct hostent {
874          char *h_name;
875          char **h_aliases;
876          int h_addrtype;
877          int h_length;
878          char **h_addr_list;
879      };
880      struct protoent {
881          char *p_name;
882          char **p_aliases;
883          int p_proto;
884      };
885      struct netent {
886          char *n_name;
887          char **n_aliases;
888          int n_addrtype;
889          unsigned int n_net;
890      };
891
892      #define AI_PASSIVE 0x0001
893      #define AI_CANONNAME 0x0002
894      #define AI_NUMERICHOST 0x0004
895
896      struct addrinfo {
897          int ai_flags;
898          int ai_family;
899          int ai_socktype;
900          int ai_protocol;
901          socklen_t ai_addrlen;
902          struct sockaddr *ai_addr;
903          char *ai_canonname;
904          struct addrinfo *ai_next;
905      };
906
907      #define NI_NUMERICHOST 1
908      #define NI_DGRAM 16
909      #define NI_NUMERICSERV 2

```

```

910 #define NI_NOFQDN      4
911 #define NI_NAMEREQD    8
912
913 #define EAI_BADFLAGS   -1
914 #define EAI_MEMORY     -10
915 #define EAI_SYSTEM     -11
916 #define EAI_NONAME    -2
917 #define EAI AGAIN     -3
918 #define EAI FAIL      -4
919 #define EAI NODATA    -5
920 #define EAI FAMILY    -6
921 #define EAI SOCKTYPE  -7
922 #define EAI SERVICE   -8
923 #define EAI ADDRFAMILY -9
924
925 extern void endprotoent(void);
926 extern void endservent(void);
927 extern void freeaddrinfo(struct addrinfo *);
928 extern const char *gai_strerror(int);
929 extern int getaddrinfo(const char *, const char *, const struct addrinfo
930 *,
931             struct addrinfo **);
932 extern struct hostent *gethostbyaddr(const void *, socklen_t, int);
933 extern struct hostent *gethostbyname(const char *);
934 extern struct protoent *getprotobynumber(const char *);
935 extern struct protoent *getprotoent(int);
936 extern struct protoent *getprotoent(void);
937 extern struct servent *getservbyname(const char *, const char *);
938 extern struct servent *getservbyport(int, const char *);
939 extern struct servent *getservent(void);
940 extern void setprotoent(int);
941 extern void setservent(int);
942 extern int *_h_errno_location(void);

```

### 13.4.24 netinet/in.h

```

943
944 #define IPPROTO_IP      0
945 #define IPPROTO_ICMP    1
946 #define IPPROTO_UDP    17
947 #define IPPROTO_IGMP    2
948 #define IPPROTO_RAW     255
949 #define IPPROTO_IPV6   41
950 #define IPPROTO_ICMPV6 58
951 #define IPPROTO_TCP    6
952
953 typedef uint16_t in_port_t;
954
955 struct in_addr {
956     uint32_t s_addr;
957 };
958 typedef uint32_t in_addr_t;
959
960 #define INADDR_NONE      ((in_addr_t) 0xffffffff)
961 #define INADDR_BROADCAST ((0xffffffff))
962 #define INADDR_ANY        0
963
964 struct in6_addr {
965     union {
966         uint8_t u6_addr8[16];
967         uint16_t u6_addr16[8];
968         uint32_t u6_addr32[4];
969     } in6_u;
970 };

```

```

971
972     #define IN6ADDR_ANY_INIT      { { { 0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0 } } }
973     #define IN6ADDR_LOOPBACK_INIT { { { 0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,1 } } }
974
975
976     #define INET_ADDRSTRLEN 16
977
978     struct sockaddr_in {
979         sa_family_t sin_family;
980         unsigned short sin_port;
981         struct in_addr sin_addr;
982         unsigned char sin_zero[8];
983     };
984
985     #define INET6_ADDRSTRLEN        46
986
987     struct sockaddr_in6 {
988         unsigned short sin6_family;
989         uint16_t sin6_port;
990         uint32_t sin6_flowinfo;
991         struct in6_addr sin6_addr;
992         uint32_t sin6_scope_id;
993     };
994
995     #define SOL_IP    0
996     #define IP_TOS   1
997     #define IPV6_UNICAST_HOPS    16
998     #define IPV6_MULTICAST_IF    17
999     #define IPV6_MULTICAST_HOPS    18
1000    #define IPV6_MULTICAST_LOOP    19
1001    #define IP_TTL   2
1002    #define IPV6_JOIN_GROUP 20
1003    #define IPV6_LEAVE_GROUP    21
1004    #define IPV6_V6ONLY    26
1005    #define IP_MULTICAST_IF    32
1006    #define IP_MULTICAST_TTL    33
1007    #define IP_MULTICAST_LOOP    34
1008    #define IP_ADD_MEMBERSHIP    35
1009    #define IP_DROP_MEMBERSHIP    36
1010    #define IP_OPTIONS    4
1011
1012    struct ipv6_mreq {
1013        struct in6_addr ipv6mr_multiaddr;
1014        int ipv6mr_interface;
1015    };
1016    struct ip_mreq {
1017        struct in_addr imr_multiaddr;
1018        struct in_addr imr_interface;
1019    };
1020    extern int bindresvport(int, struct sockaddr_in *);

```

### 13.4.25 netinet/ip.h

```

1021
1022     #define IPTOS_LOWCOST    0x02
1023     #define IPTOS_RELIABILITY 0x04
1024     #define IPTOS_THROUGHPUT 0x08
1025     #define IPTOS_LOWDELAY   0x10
1026     #define IPTOS_TOS_MASK   0x1e
1027     #define IPTOS_MINCOST    IPTOS_LOWCOST
1028
1029     #define IPTOS_PREC_MASK  0xe0

```

### 13.4.26 netinet/tcp.h

```
1030
1031     #define TCP_NODELAY      1
1032     #define SOL_TCP 6
```

### 13.4.27 netinet/udp.h

```
1033
1034     #define SOL_UDP 17
```

### 13.4.28 nl\_types.h

```
1035
1036     #define NL_CAT_LOCALE    1
1037     #define NL_SETD 1
1038
1039     typedef void *nl_catd;
1040
1041     typedef int nl_item;
1042     extern int catclose(nl_catd);
1043     extern char *catgets(nl_catd, int, int, const char *);
1044     extern nl_catd catopen(const char *, int);
```

### 13.4.29 poll.h

```
1045
1046     extern int poll(struct pollfd *, nfds_t, int);
```

### 13.4.30 pty.h

```
1047
1048     extern int openpty(int *, int *, char *, struct termios *,
1049                         struct winsize *);
1050     extern int forkpty(int *, char *, struct termios *, struct winsize *);
```

### 13.4.31 pwd.h

```
1051
1052     struct passwd {
1053         char *pw_name;
1054         char *pw_passwd;
1055         uid_t pw_uid;
1056         gid_t pw_gid;
1057         char *pw_gecos;
1058         char *pw_dir;
1059         char *pw_shell;
1060     };
1061     extern void endpwent(void);
1062     extern struct passwd *getpwent(void);
1063     extern struct passwd *getpwnam(char *);
1064     extern struct passwd *getpwuid(uid_t);
1065     extern void setpwent(void);
1066     extern int getpwnam_r(char *, struct passwd *, char *, size_t,
1067                           struct passwd **);
1068     extern int getpwuid_r(uid_t, struct passwd *, char *, size_t,
1069                           struct passwd **);
```

### 13.4.32 regex.h

```
1070
1071     typedef unsigned long int reg_syntax_t;
```

```

1072
1073     typedef struct re_pattern_buffer {
1074         unsigned char *buffer;
1075         unsigned long int allocated;
1076         unsigned long int used;
1077         reg_syntax_t syntax;
1078         char *fastmap;
1079         char *translate;
1080         size_t re_nsub;
1081         unsigned int can_be_null:1;
1082         unsigned int regs_allocated:2;
1083         unsigned int fastmap_accurate:1;
1084         unsigned int no_sub:1;
1085         unsigned int not_bol:1;
1086         unsigned int not_eol:1;
1087         unsigned int newline_anchor:1;
1088     } regex_t;
1089     typedef int regoff_t;
1090     typedef struct {
1091         regoff_t rm_so;
1092         regoff_t rm_eo;
1093     } regmatch_t;
1094
1095 #define REG_ICASE      (REG_EXTENDED<<1)
1096 #define REG_NEWLINE    (REG_ICASE<<1)
1097 #define REG_NOSUB      (REG_NEWLINE<<1)
1098 #define REG_EXTENDED   1
1099
1100 #define REG_NOTEOL     (1<<1)
1101 #define REG_NOTBOL     1
1102
1103     typedef enum {
1104         REG_ENOSYS, REG_NOERROR, REG_NOMATCH, REG_BADPAT, REG_ECOLLATE,
1105             REG_ECTYPE, REG_EESCAPE, REG_ESUBREG, REG_EBRACK, REG_EPAREN,
1106             REG_EBRACE, REG_BADBR, REG_ERANGE, REG_ESPACE, REG_BADRPT,
1107             REG_EEND, REG_ESIZE, REG_ERPAREN
1108     } reg_errcode_t;
1109     extern int regcomp(regex_t *, const char *, int);
1110     extern size_t regerror(int, const regex_t *, char *, size_t);
1111     extern int regexec(const regex_t *, const char *, size_t, regmatch_t,
1112     int);
1113     extern void regfree(regex_t *);

```

### 13.4.33 rpc/auth.h

```

1114
1115     enum auth_stat {
1116         AUTH_OK, AUTH_BADCRED = 1, AUTH_REJECTEDCRED = 2, AUTH_BADVERF =
1117             3, AUTH_REJECTEDVERF = 4, AUTH_TOOWEAK = 5, AUTH_INVALIDRESP =
1118             6, AUTH_FAILED = 7
1119     };
1120
1121     union des_block {
1122         struct {
1123             u_int32_t high;
1124             u_int32_t low;
1125         } key;
1126         char c[8];
1127     };
1128
1129     struct opaque_auth {
1130         enum_t oa_flavor;
1131         caddr_t oa_base;
1132         u_int oa_length;

```

```

1133 } ;
1134
1135 typedef struct AUTH {
1136     struct opaque_auth ah_cred;
1137     struct opaque_auth ah_verf;
1138     union des_block ah_key;
1139     struct auth_ops *ah_ops;
1140     caddr_t ah_private;
1141 } AUTH;
1142
1143 struct auth_ops {
1144     void (*ah_nextverf) (struct AUTH *);
1145     int (*ah_marshall) (struct AUTH *, XDR *);
1146     int (*ah_validate) (struct AUTH *, struct opaque_auth *);
1147     int (*ah_refresh) (struct AUTH *);
1148     void (*ah_destroy) (struct AUTH *);
1149 };
1150 extern struct AUTH *authnone_create(void);
1151 extern int key_decryptsession(char *, union des_block *);
1152 extern bool_t xdr_opaque_auth(XDR *, struct opaque_auth *);

```

### 13.4.34 rpc/clnt.h

```

1153 #define clnt_control(cl,rq,in)
1154     (((cl)->cl_ops->cl_control)(cl,rq,in))
1155 #define clnt_abort(rh)    (((rh)->cl_ops->cl_abort)(rh))
1156 #define clnt_destroy(rh)   (((rh)->cl_ops->cl_destroy)(rh))
1157 #define clnt_freeres(rh,xres,resp)
1158     (((rh)->cl_ops->cl_freeres)(rh,xres,resp))
1159 #define clnt_geterr(rh,errp)    (((rh)->cl_ops->cl_geterr)(rh, errp))
1160 #define NULLPROC          ((u_long)0)
1161 #define CLSET_TIMEOUT      1
1162 #define CLGET_XID          10
1163 #define CLSET_XID          11
1164 #define CLGET_VERS          12
1165 #define CLSET_VERS          13
1166 #define CLGET_PROG          14
1167 #define CLSET_PROG          15
1168 #define CLGET_TIMEOUT       2
1169 #define CLGET_SERVER_ADDR    3
1170 #define CLSET_RETRY_TIMEOUT  4
1171 #define CLGET_RETRY_TIMEOUT  5
1172 #define CLGET_FD             6
1173 #define CLGET_SVC_ADDR       7
1174 #define CLSET_FD_CLOSE        8
1175 #define CLSET_FD_NCLOSE      9
1176 #define clnt_call(rh, proc, xargs, argsp, xres, resp, secs) \
1177     (((rh)->cl_ops->cl_call)(rh, proc, xargs, argsp, xres, resp, \
1178     secs))
1179
1180 enum clnt_stat {
1181     RPC_SUCCESS, RPC_CANTENCODEARGS = 1, RPC_CANTDECODERES =
1182         2, RPC_CANTSEND = 3, RPC_CANTRECV = 4, RPC_TIMEDOUT =
1183         5, RPC_VERSMISMATCH = 6, RPC_AUTHERROR = 7, RPC_PROGUNAVAIL =
1184         8, RPC_PROGVERSISMISMATCH = 9, RPC_PROCUNAVAIL =
1185         10, RPC_CANTDECODEARGS = 11, RPC_SYSTEMERROR =
1186         12, RPC_NOBROADCAST = 21, RPC_UNKNOWNHOST = 13, RPC_UNKNOWNPROTO
1187     =
1188         17, RPC_UNKNOWNADDR = 19, RPC_RPCBFAILURE =
1189         14, RPC_PROGNOTREGISTERED = 15, RPC_N2AXLATEFAILURE =
1190         22, RPC_FAILED = 16, RPC_INTR = 18, RPC_TLIERROR =
1191         20, RPC_UDERROR = 23, RPC_INPROGRESS = 24, RPC_STALERACHANDLE
1192     =
1193         25

```

```

1194     };
1195     struct rpc_err {
1196         enum clnt_stat re_status;
1197         union {
1198             int RE_errno;
1199             enum auth_stat RE_why;
1200             struct {
1201                 u_long low;
1202                 u_long high;
1203             } RE_vers;
1204             struct {
1205                 long int s1;
1206                 long int s2;
1207             } RE_lb;
1208         } ru;
1209     };
1210
1211     typedef struct CLIENT {
1212         struct AUTH *cl_auth;
1213         struct clnt_ops *cl_ops;
1214         caddr_t cl_private;
1215     } CLIENT;
1216
1217     struct clnt_ops {
1218         enum clnt_stat (*cl_call) (struct CLIENT *, u_long, xdrproc_t,
1219                                   caddr_t,
1220                                   xdrproc_t, caddr_t, struct timeval);
1221         void (*cl_abort) (void);
1222         void (*cl_geterr) (struct CLIENT *, struct rpc_err *);
1223         bool_t (*cl_freeres) (struct CLIENT *, xdrproc_t, caddr_t);
1224         void (*cl_destroy) (struct CLIENT *);
1225         bool_t (*cl_control) (struct CLIENT *, int, char *);
1226     };
1227     extern struct CLIENT *clnt_create(const char *, const u_long, const
1228                                       u_long,
1229                                       const char *);
1230     extern void clnt_pcreateerror(const char *);
1231     extern void clnt_perrno(enum clnt_stat);
1232     extern void clnt_perror(struct CLIENT *, const char *);
1233     extern char *clnt_spcreateerror(const char *);
1234     extern char *clnt_sperrno(enum clnt_stat);
1235     extern char *clnt_sperror(struct CLIENT *, const char );

```

### 13.4.35 rpc/pmap\_clnt.h

```

1236     extern u_short pmap_getport(struct sockaddr_in *, const u_long,
1237                                   const u_long, u_int);
1238     extern bool_t pmap_set(const u_long, const u_long, int, u_short);
1239     extern bool_t pmap_unset(u_long, u_long);
1240

```

### 13.4.36 rpc/rpc\_msg.h

```

1241     enum msg_type {
1242         CALL, REPLY = 1
1243     };
1244     enum reply_stat {
1245         MSG_ACCEPTED, MSG_DENIED = 1
1246     };
1247     enum accept_stat {
1248         SUCCESS, PROG_UNAVAIL = 1, PROG_MISMATCH = 2, PROC_UNAVAIL =
1249             3, GARBAGE_ARGS = 4, SYSTEM_ERR = 5
1250     };
1251

```

```

1252     enum reject_stat {
1253         RPC_MISMATCH, AUTH_ERROR = 1
1254     };
1255
1256     struct accepted_reply {
1257         struct opaque_auth ar_verf;
1258         enum accept_stat ar_stat;
1259         union {
1260             struct {
1261                 unsigned long int low;
1262                 unsigned long int high;
1263             } AR_versions;
1264             struct {
1265                 caddr_t where;
1266                 xdrproc_t proc;
1267             } AR_results;
1268         } ru;
1269     };
1270
1271     struct rejected_reply {
1272         enum reject_stat rj_stat;
1273         union {
1274             struct {
1275                 unsigned long int low;
1276                 unsigned long int high;
1277             } RJ_versions;
1278             enum auth_stat RJ_why;
1279         } ru;
1280     };
1281
1282     struct reply_body {
1283         enum reply_stat rp_stat;
1284         union {
1285             struct accepted_reply RP_ar;
1286             struct rejected_reply RP_dr;
1287         } ru;
1288     };
1289
1290     struct call_body {
1291         unsigned long int cb_rpcvers;
1292         unsigned long int cb_prog;
1293         unsigned long int cb_vers;
1294         unsigned long int cb_proc;
1295         struct opaque_auth cb_cred;
1296         struct opaque_auth cb_verf;
1297     };
1298
1299     struct rpc_msg {
1300         unsigned long int rm_xid;
1301         enum msg_type rm_direction;
1302         union {
1303             struct call_body RM_cmb;
1304             struct reply_body RM_rmb;
1305         } ru;
1306     };
1307     extern bool_t xdr_callhdr(XDR *, struct rpc_msg *);

```

### 13.4.37 rpc/svc.h

```

1308
1309 #define RPC_ANYSOCK      -1
1310 #define svc_freeargs(xprt,xargs, argsp) \
1311     (*(xprt)->xp_ops->xp_freeargs)((xprt), (xargs), (argsp))
1312 #define svc_getargs(xprt,xargs, argsp) \

```

```

1313         (*(xpert)->xp_ops->xp_getargs)((xpert), (xargs), (argsp))
1314
1315     enum xpert_stat {
1316         XPRT_DIED, XPRT_MOREREQS, XPRT_IDLE
1317     };
1318
1319     typedef struct SVCXPRT {
1320         int xp_sock;
1321         u_short xp_port;
1322         struct xp_ops *xp_ops;
1323         int xp_addrlen;
1324         struct sockaddr_in xp_raddr;
1325         struct opaque_auth xp_verf;
1326         caddr_t xp_p1;
1327         caddr_t xp_p2;
1328         char xp_pad[256];
1329     } SVCXPRT;
1330
1331     struct svc_req {
1332         rpcprog_t rq_prog;
1333         rpcvers_t rq_vers;
1334         rpcproc_t rq_proc;
1335         struct opaque_auth rq_cred;
1336         caddr_t rq_clntcred;
1337         SVCXPRT *rq_xpert;
1338     };
1339
1340     typedef void (*__dispatch_fn_t) (struct svc_req *, SVCXPRT *);
1341
1342     struct xp_ops {
1343         bool_t(*xp_recv) (SVCXPRT * __xpert, struct rpc_msg * __msg);
1344         enum xpert_stat (*xp_stat) (SVCXPRT * __xpert);
1345         bool_t(*xp_getargs) (SVCXPRT * __xpert, xdrproc_t __xdr_args,
1346                             caddr_t args_ptr);
1347         bool_t(*xp_reply) (SVCXPRT * __xpert, struct rpc_msg * __msg);
1348         bool_t(*xp_freeargs) (SVCXPRT * __xpert, xdrproc_t __xdr_args,
1349                               caddr_t args_ptr);
1350         void (*xp_destroy) (SVCXPRT * __xpert);
1351     };
1352     extern void svc_getreqset(fd_set *);
1353     extern bool_t svc_register(SVCXPRT *, rpcprog_t, rpcvers_t,
1354                                __dispatch_fn_t, rpcprot_t);
1355     extern void svc_run(void);
1356     extern bool_t svc_sendreply(SVCXPRT *, xdrproc_t, caddr_t);
1357     extern void svcerr_auth(SVCXPRT *, enum auth_stat);
1358     extern void svcerr_decode(SVCXPRT *);
1359     extern void svcerr_noproc(SVCXPRT *);
1360     extern void svcerr_noprog(SVCXPRT *);
1361     extern void svcerr_progvers(SVCXPRT *, rpcvers_t, rpcvers_t);
1362     extern void svcerr_systemerr(SVCXPRT *);
1363     extern void svcerr_weakauth(SVCXPRT *);
1364     extern SVCXPRT *svctcp_create(int, u_int, u_int);
1365     extern SVCXPRT *svcudp_create(int);

```

### 13.4.38 rpc/types.h

```

1366
1367     typedef int bool_t;
1368     typedef int enum_t;
1369     typedef unsigned long int rpcprog_t;
1370     typedef unsigned long int rpcvers_t;
1371     typedef unsigned long int rpcproc_t;
1372     typedef unsigned long int rpcprot_t;

```

### 13.4.39 rpc/xdr.h

```

1373
1374     enum xdr_op {
1375         XDR_ENCODE, XDR_DECODE, XDR_FREE
1376     };
1377     typedef struct XDR {
1378         enum xdr_op x_op;
1379         struct xdr_ops *x_ops;
1380         caddr_t x_public;
1381         caddr_t x_private;
1382         caddr_t x_base;
1383         int x_handy;
1384     } XDR;
1385
1386     struct xdr_ops {
1387         bool_t(*x_getlong) (XDR * __xdrs, long int *__lp);
1388         bool_t(*x_putlong) (XDR * __xdrs, long int *__lp);
1389         bool_t(*x_getbytes) (XDR * __xdrs, caddr_t __addr, u_int __len);
1390         bool_t(*x_putbytes) (XDR * __xdrs, char *__addr, u_int __len);
1391         u_int(*x_getpostn) (XDR * __xdrs);
1392         bool_t(*x_setpostn) (XDR * __xdrs, u_int __pos);
1393         int32_t *(*x_inline) (XDR * __xdrs, int __len);
1394         void (*x_destroy) (XDR * __xdrs);
1395         bool_t(*x_getint32) (XDR * __xdrs, int32_t * __ip);
1396         bool_t(*x_putint32) (XDR * __xdrs, int32_t * __ip);
1397     };
1398
1399     typedef bool_t(*xdrproc_t) (XDR *, void *, ...);
1400
1401     struct xdr_discrim {
1402         int value;
1403         xdrproc_t proc;
1404     };
1405     extern bool_t xdr_array(XDR *, caddr_t *, u_int *, u_int, u_int,
1406                             xdrproc_t);
1407     extern bool_t xdr_bool(XDR *, bool_t *);
1408     extern bool_t xdr_bytes(XDR *, char **, u_int *, u_int);
1409     extern bool_t xdr_char(XDR *, char *);
1410     extern bool_t xdr_double(XDR *, double *);
1411     extern bool_t xdr_enum(XDR *, enum_t *);
1412     extern bool_t xdr_float(XDR *, float *);
1413     extern void xdr_free(xdrproc_t, char *);
1414     extern bool_t xdr_int(XDR *, int *);
1415     extern bool_t xdr_long(XDR *, long int *);
1416     extern bool_t xdr_opaque(XDR *, caddr_t, u_int);
1417     extern bool_t xdr_pointer(XDR *, char **, u_int, xdrproc_t);
1418     extern bool_t xdr_reference(XDR *, caddr_t *, u_int, xdrproc_t);
1419     extern bool_t xdr_short(XDR *, short *);
1420     extern bool_t xdr_string(XDR *, char **, u_int);
1421     extern bool_t xdr_u_char(XDR *, u_char *);
1422     extern bool_t xdr_u_int(XDR *, u_int *);
1423     extern bool_t xdr_u_long(XDR *, u_long *);
1424     extern bool_t xdr_u_short(XDR *, u_short *);
1425     extern bool_t xdr_union(XDR *, enum_t *, char *,
1426                            const struct xdr_discrim *, xdrproc_t);
1427     extern bool_t xdr_vector(XDR *, char *, u_int, u_int, xdrproc_t);
1428     extern bool_t xdr_void(void);
1429     extern bool_t xdr_wrapstring(XDR *, char **);
1430     extern void xdrmem_create(XDR *, caddr_t, u_int, enum xdr_op);
1431     extern void xdrrec_create(XDR *, u_int, u_int, caddr_t,
1432                              int (*__readit) (char *p1, char *p2, int p3)
1433                              , int (*__writeit) (char *p1, char *p2, int
1434 p3))

```

```
1435     );
1436     extern typedef int bool_t xdrrec_eof(XDR *);
```

### 13.4.40 sched.h

```
1437
1438 #define SCHED_OTHER      0
1439 #define SCHED_FIFO       1
1440 #define SCHED_RR        2
1441
1442     struct sched_param {
1443         int sched_priority;
1444     };
1445     extern int sched_get_priority_max(int);
1446     extern int sched_get_priority_min(int);
1447     extern int sched_getparam(pid_t, struct sched_param *);
1448     extern int sched_setscheduler(pid_t);
1449     extern int sched_rr_get_interval(pid_t, struct timespec *);
1450     extern int sched_setparam(pid_t, const struct sched_param *);
1451     extern int sched_setscheduler(pid_t, int, const struct sched_param *);
1452     extern int sched_yield(void);
```

### 13.4.41 search.h

```
1453
1454     typedef struct entry {
1455         char *key;
1456         void *data;
1457     } ENTRY;
1458     typedef enum {
1459         FIND, ENTER
1460     } ACTION;
1461     typedef enum {
1462         preorder, postorder, endorder, leaf
1463     } VISIT;
1464
1465     typedef void (*__action_fn_t) (void *__nodep, VISIT __value, int
1466     __level);
1467     extern int hcreate(size_t);
1468     extern ENTRY *hsearch(ENTRY, ACTION);
1469     extern void insque(void *, void *);
1470     extern void *lfind(const void *, const void *, size_t *, size_t,
1471     __compar_fn_t);
1472     extern void *lsearch(const void *, void *, size_t *, size_t,
1473     __compar_fn_t);
1474     extern void remque(void *);
1475     extern void hdestroy(void);
1476     extern void *tdelete(const void *, void **, __compar_fn_t);
1477     extern void *tfind(const void *, void *const *, __compar_fn_t);
1478     extern void *tsearch(const void *, void **, __compar_fn_t);
1479     extern void twalk(const void *, __action_fn_t);
```

### 13.4.42 setjmp.h

```
1480
1481 #define setjmp(env)      _setjmp(env)
1482 #define sigsetjmp(a,b)   __sigsetjmp(a,b)
1483
1484     struct __jmp_buf_tag {
1485         __jmp_buf __jmpbuf;
1486         int __mask_was_saved;
1487         sigset_t __saved_mask;
1488     };
```

```

1489
1490     typedef struct __jmp_buf_tag jmp_buf[1];
1491     typedef jmp_buf sigjmp_buf;
1492     extern int __sigsetjmp(jmp_buf, int);
1493     extern void longjmp(jmp_buf, int);
1494     extern void siglongjmp(sigjmp_buf, int);
1495     extern void _longjmp(jmp_buf, int);
1496     extern int _setjmp(jmp_buf);

```

### 13.4.43 signal.h

```

1497
1498 #define _SIGSET_NWORDS  (1024/(8*sizeof(unsigned long)))
1499 #define SIGRTMAX      (__libc_current_sigrtmax ())
1500 #define SIGRTMIN      (__libc_current_sigrtmin ())
1501 #define SIG_BLOCK      0
1502 #define SIG_UNBLOCK    1
1503 #define SIG_SETMASK    2
1504 #define NSIG          65
1505
1506     typedef int sig_atomic_t;
1507
1508     typedef void (*sighandler_t) (int);
1509
1510 #define SIG_HOLD        ((sighandler_t) 2)
1511 #define SIG_ERR         ((sighandler_t)-1)
1512 #define SIG_DFL         ((sighandler_t)0)
1513 #define SIG_IGN         ((sighandler_t)1)
1514
1515 #define SIGHUP          1
1516 #define SIGUSR1         10
1517 #define SIGSEGV         11
1518 #define SIGUSR2         12
1519 #define SIGPIPE         13
1520 #define SIGALRM         14
1521 #define SIGTERM         15
1522 #define SIGSTKFLT       16
1523 #define SIGCHLD         17
1524 #define SIGCONT         18
1525 #define SIGSTOP         19
1526 #define SIGINT          2
1527 #define SIGTSTP         20
1528 #define SIGTTIN         21
1529 #define SIGTTOU         22
1530 #define SIGURG          23
1531 #define SIGXCPU         24
1532 #define SIGXFSZ         25
1533 #define SIGVTALRM       26
1534 #define SIGPROF         27
1535 #define SIGWINCH        28
1536 #define SIGIO           29
1537 #define SIGQUIT         3
1538 #define SIGPWR          30
1539 #define SIGSYS          31
1540 #define SIGUNUSED        31
1541 #define SIGILL          4
1542 #define SIGTRAP         5
1543 #define SIGABRT         6
1544 #define SIGIOT          6
1545 #define SIGBUS          7
1546 #define SIGFPE          8
1547 #define SIGKILL         9
1548 #define SIGCLD          SIGCHLD
1549 #define SIGPOLL         SIGIO

```

```

1550
1551     #define SV_ONSTACK      (1<<0)
1552     #define SV_INTERRUPT    (1<<1)
1553     #define SV_RESETHAND   (1<<2)
1554
1555     typedef union sigval {
1556         int sival_int;
1557         void *sival_ptr;
1558     } sigval_t;
1559
1560     #define SIGEV_SIGNAL     0
1561     #define SIGEV_NONE       1
1562     #define SIGEV_THREAD     2
1563     #define SIGEV_MAX_SIZE   64
1564
1565     typedef struct sigevent {
1566         sigval_t sigev_value;
1567         int sigev_signo;
1568         int sigev_notify;
1569         union {
1570             int _pad[SIGEV_PAD_SIZE];
1571             struct {
1572                 void (*sigev_thread_func) (sigval_t);
1573                 void *_attribute;
1574             } _sigev_thread;
1575         } _sigev_un;
1576     } sigevent_t;
1577
1578     #define SI_MAX_SIZE      128
1579     #define si_pid  _sifields._kill._pid
1580     #define si_uid  _sifields._kill._uid
1581     #define si_value _sifields._rt._sigval
1582     #define si_int   _sifields._rt._sigval.sival_int
1583     #define si_ptr   _sifields._rt._sigval.sival_ptr
1584     #define si_status _sifields._sigchld._status
1585     #define si_stime _sifields._sigchld._stime
1586     #define si_utime _sifields._sigchld._utime
1587     #define si_addr _sifields._sigfault._addr
1588     #define si_band _sifields._sigpoll._band
1589     #define si_fd   _sifields._sigpoll._fd
1590     #define si_timer1 _sifields._timer._timer1
1591     #define si_timer2 _sifields._timer._timer2
1592
1593     typedef struct siginfo {
1594         int si_signo;
1595         int si_errno;
1596         int si_code;
1597         union {
1598             int _pad[SI_PAD_SIZE];
1599             struct {
1600                 pid_t _pid;
1601                 uid_t _uid;
1602             } _kill;
1603             struct {
1604                 unsigned int _timer1;
1605                 unsigned int _timer2;
1606             } _timer;
1607             struct {
1608                 pid_t _pid;
1609                 uid_t _uid;
1610                 sigval_t _sigval;
1611             } _rt;
1612             struct {
1613                 pid_t _pid;

```

```

1614             uid_t _uid;
1615             int _status;
1616             clock_t _utime;
1617             clock_t _stime;
1618         } _sigchld;
1619         struct {
1620             void *_addr;
1621         } _sigfault;
1622         struct {
1623             int _band;
1624             int _fd;
1625         } _sigpoll;
1626     } _sifields;
1627 } siginfo_t;
1628
1629 #define SI_QUEUE      -1
1630 #define SI_TIMER      -2
1631 #define SI_MESGQ      -3
1632 #define SI_ASYNCIO    -4
1633 #define SI_SIGIO      -5
1634 #define SI_TKILL      -6
1635 #define SI_ASYNCNL    -60
1636 #define SI_USER 0
1637 #define SI_KERNEL      0x80
1638
1639 #define ILL_ILLOPC    1
1640 #define ILL_ILLOPN    2
1641 #define ILL_ILLADR    3
1642 #define ILL_ILLTRP    4
1643 #define ILL_PRVOPC   5
1644 #define ILL_PRVREG   6
1645 #define ILL_COPROC    7
1646 #define ILL_BADSTK   8
1647
1648 #define FPE_INTDIV   1
1649 #define FPE_INTOVF   2
1650 #define FPE_FLTDIV   3
1651 #define FPE_FLTOVF   4
1652 #define FPE_FLTUND   5
1653 #define FPE_FLTRES   6
1654 #define FPE_FLTINV   7
1655 #define FPE_FLTSUB   8
1656
1657 #define SEGV_MAPERR   1
1658 #define SEGV_ACCERR   2
1659
1660 #define BUS_ADRALN   1
1661 #define BUS adrerr   2
1662 #define BUS_OBJERR   3
1663
1664 #define TRAP_BRKPT   1
1665 #define TRAP_TRACE    2
1666
1667 #define CLD_EXITED   1
1668 #define CLD_KILLED   2
1669 #define CLD_DUMPED   3
1670 #define CLD_TRAPPED   4
1671 #define CLD_STOPPED   5
1672 #define CLD_CONTINUED 6
1673
1674 #define POLL_IN 1
1675 #define POLL_OUT 2
1676 #define POLL_MSG 3
1677 #define POLL_ERR 4

```

```

1678     #define POLL_PRI      5
1679     #define POLL_HUP       6
1680
1681     typedef struct {
1682         unsigned long int sig[_SIGSET_NWORDS];
1683     } sigset_t;
1684
1685     #define SA_NOCLDSTOP   0x00000001
1686     #define SA_NOCLDWAIT   0x00000002
1687     #define SA_SIGINFO      0x00000004
1688     #define SA_ONSTACK      0x08000000
1689     #define SA_RESTART      0x10000000
1690     #define SA_INTERRUPT    0x20000000
1691     #define SA_NODEFER      0x40000000
1692     #define SA_RESETHAND   0x80000000
1693     #define SA_NOMASK      SA_NODEFER
1694     #define SA_ONESHOT     SA_RESETHAND
1695
1696     typedef struct sigaltstack {
1697         void *ss_sp;
1698         int ss_flags;
1699         size_t ss_size;
1700     } stack_t;
1701
1702     #define SS_ONSTACK     1
1703     #define SS_DISABLE      2
1704
1705     extern int __libc_current_sigrtmax(void);
1706     extern int __libc_current_sigrtmin(void);
1707     extern sighandler_t __sysv_signal(int, sighandler_t);
1708     extern char *const _sys_siglist(void);
1709     extern int killpg(pid_t, int);
1710     extern void psignal(int, const char *);
1711     extern int raise(int);
1712     extern int sigaddset(sigset_t *, int);
1713     extern int sigandset(sigset_t *, const sigset_t *, const sigset_t *);
1714     extern int sigdelset(sigset_t *, int);
1715     extern int sigemptyset(sigset_t *);
1716     extern int sigfillset(sigset_t *);
1717     extern int sighold(int);
1718     extern int sigignore(int);
1719     extern int siginterrupt(int, int);
1720     extern int sigisemptyset(const sigset_t *);
1721     extern int sigismember(const sigset_t *, int);
1722     extern int sigorset(sigset_t *, const sigset_t *, const sigset_t *);
1723     extern int sigpending(sigset_t *);
1724     extern int sigrelse(int);
1725     extern sighandler_t sigset(int, sighandler_t);
1726     extern int pthread_kill(pthread_t, int);
1727     extern int pthread_sigmask(int, sigset_t *, sigset_t *);
1728     extern int sigaction(int, const struct sigaction *, struct sigaction *);
1729     extern int sigwait(sigset_t *, int *);
1730     extern int kill(pid_t, int);
1731     extern int sigaltstack(const struct sigaltstack *, struct sigaltstack *);
1732
1733     extern sighandler_t signal(int, sighandler_t);
1734     extern int sigpause(int);
1735     extern int sigprocmask(int, const sigset_t *, sigset_t *);
1736     extern int sigreturn(struct sigcontext *);
1737     extern int sigsuspend(const sigset_t *);
1738     extern int sigqueue(pid_t, int, const union sigval);
1739     extern int sigwaitinfo(const sigset_t *, siginfo_t *);
1740     extern int sigtimedwait(const sigset_t *, siginfo_t *,
1741                           const struct timespec *);

```

```
1742     extern sighandler_t bsd_signal(int, sighandler_t);
```

### 13.4.44 stddef.h

```
1743
1744 #define offsetof(TYPE, MEMBER)    ((size_t)&((TYPE*)0)->MEMBER)
1745 #define NULL          (0L)
1746
1747 typedef int wchar_t;
```

### 13.4.45 stdio.h

```
1748
1749 #define EOF      (-1)
1750 #define P_tmpdir      "/tmp"
1751 #define FOPEN_MAX      16
1752 #define L_tmpnam      20
1753 #define FILENAME_MAX  4096
1754 #define BUFSIZ      8192
1755 #define L_ctermid     9
1756 #define L_cuserid      9
1757
1758 typedef struct {
1759     off_t __pos;
1760     mbstate_t __state;
1761 } fpos_t;
1762 typedef struct {
1763     off64_t __pos;
1764     mbstate_t __state;
1765 } fpos64_t;
1766
1767 typedef struct _IO_FILE FILE;
1768
1769 #define _IOFBF    0
1770 #define _IOLBF    1
1771 #define _IONBF    2
1772
1773 extern char *const _sys_errlist(void);
1774 extern void clearerr(FILE *);
1775 extern int fclose(FILE *);
1776 extern FILE *fdopen(int, const char *);
1777 extern int fflush_unlocked(FILE *);
1778 extern int fileno(FILE *);
1779 extern FILE *fopen(const char *, const char *);
1780 extern int fprintf(FILE *, const char *, ...);
1781 extern int fputc(int, FILE *);
1782 extern FILE *freopen(const char *, const char *, FILE *);
1783 extern FILE *freopen64(const char *, const char *, FILE *);
1784 extern int fscanf(FILE *, const char *, ...);
1785 extern int fseek(FILE *, long int, int);
1786 extern int fseeko(FILE *, off_t, int);
1787 extern int fseeko64(FILE *, loff_t, int);
1788 extern off_t ftello(FILE *);
1789 extern loff_t ftello64(FILE *);
1790 extern int getchar(void);
1791 extern int getchar_unlocked(void);
1792 extern int getw(FILE *);
1793 extern int pclose(FILE *);
1794 extern void perror(const char *);
1795 extern FILE *popen(const char *, const char *);
1796 extern int printf(const char *, ...);
1797 extern int putc_unlocked(int, FILE *);
1798 extern int putchar(int);
1799 extern int putchar_unlocked(int);
```

```

1800     extern int putw(int, FILE *);
1801     extern int remove(const char *);
1802     extern void rewind(FILE *);
1803     extern int scanf(const char *, ...);
1804     extern void setbuf(FILE *, char *);
1805     extern int sprintf(char *, const char *, ...);
1806     extern int sscanf(const char *, const char *, ...);
1807     extern FILE *stderr(void);
1808     extern FILE *stdin(void);
1809     extern FILE *stdout(void);
1810     extern char *tempnam(const char *, const char *);
1811     extern FILE *tmpfile64(void);
1812     extern FILE *tmpfile(void);
1813     extern char *tmpnam(char *);
1814     extern int vfprintf(FILE *, const char *, va_list);
1815     extern int vprintf(const char *, va_list);
1816     extern int feof(FILE *);
1817     extern int ferror(FILE *);
1818     extern int fflush(FILE *);
1819     extern int fgetc(FILE *);
1820     extern int fgetpos(FILE *, fpos_t *);
1821     extern char *fgets(char *, int, FILE *);
1822     extern int fputs(const char *, FILE *);
1823     extern size_t fread(void *, size_t, size_t, FILE *);
1824     extern int fsetpos(FILE *, const fpos_t *);
1825     extern long int ftell(FILE *);
1826     extern size_t fwrite(const void *, size_t, size_t, FILE *);
1827     extern int getc(FILE *);
1828     extern int putc(int, FILE *);
1829     extern int puts(const char *);
1830     extern int setvbuf(FILE *, char *, int, size_t);
1831     extern int snprintf(char *, size_t, const char *, ...);
1832     extern int ungetc(int, FILE *);
1833     extern int vsnprintf(char *, size_t, const char *, va_list);
1834     extern int vsprintf(char *, const char *, va_list);
1835     extern void flockfile(FILE *);
1836     extern int asprintf(char **, const char *, ...);
1837     extern int fgetpos64(FILE *, fpos64_t *);
1838     extern FILE *fopen64(const char *, const char *);
1839     extern int fsetpos64(FILE *, const fpos64_t *);
1840     extern int ftrylockfile(FILE *);
1841     extern void funlockfile(FILE *);
1842     extern int getc_unlocked(FILE *);
1843     extern void setbuffer(FILE *, char *, size_t);
1844     extern int vasprintf(char **, const char *, va_list);
1845     extern int vdprintf(int, const char *, va_list);
1846     extern int vfscanf(FILE *, const char *, va_list);
1847     extern int vscanf(const char *, va_list);
1848     extern int vsscanf(const char *, const char *, va_list);
1849     extern size_t __fpending(FILE *);

```

### 13.4.46 stdlib.h

```

1850
1851 #define MB_CUR_MAX      (__ctype_get_mb_cur_max())
1852 #define EXIT_SUCCESS    0
1853 #define EXIT_FAILURE    1
1854 #define RAND_MAX        2147483647
1855
1856 typedef int (*__compar_fn_t) (const void *, const void *);
1857 struct random_data {
1858     int32_t *fptr;
1859     int32_t *rptr;
1860     int32_t *state;

```

```

1861         int rand_type;
1862         int rand_deg;
1863         int rand_sep;
1864         int32_t *end_ptr;
1865     };
1866
1867     typedef struct {
1868         int quot;
1869         int rem;
1870     } div_t;
1871
1872     typedef struct {
1873         long int quot;
1874         long int rem;
1875     } ldiv_t;
1876
1877     typedef struct {
1878         long long int quot;
1879         long long int rem;
1880     } lldiv_t;
1881
1882     extern double __strtod_internal(const char *, char **, int);
1883     extern float __strtof_internal(const char *, char **, int);
1884     extern long int __ strtol_internal(const char *, char **, int, int);
1885     extern long double __ strtold_internal(const char *, char **, int);
1886     extern long long int __ strtoll_internal(const char *, char **, int, int);
1887     extern unsigned long int __ strtoul_internal(const char *, char **, int,
1888                                                 int);
1889     extern unsigned long long int __ strtoull_internal(const char *, char **,
1890                                                       int, int);
1890     extern long int a64l(const char *);
1891     extern void abort(void);
1892     extern int abs(int);
1893     extern double atof(const char *);
1894     extern int atoi(char *);
1895     extern long int atol(char *);
1896     extern long long int atoll(const char *);
1897     extern void *bsearch(const void *, const void *, size_t, size_t,
1898                          __compar_fn_t);
1899     extern div_t div(int, int);
1900     extern double drand48(void);
1901     extern char *ecvt(double, int, int *, int *);
1902     extern double erand48(unsigned short);
1903     extern void exit(int);
1904     extern char *fcvt(double, int, int *, int *);
1905     extern char *gcvt(double, int, char *);
1906     extern char *getenv(const char *);
1907     extern int getsubopt(char **, char *const *, char **);
1908     extern int grantpt(int);
1909     extern long int jrand48(unsigned short);
1910     extern char *l64a(long int);
1911     extern long int labs(long int);
1912     extern void lcong48(unsigned short);
1913     extern ldiv_t ldiv(long int, long int);
1914     extern long long int llabs(long long int);
1915     extern lldiv_t lldiv(long long int, long long int);
1916     extern long int lrand48(void);
1917     extern int mblen(const char *, size_t);
1918     extern size_t mbstowcs(wchar_t *, const char *, size_t);
1919     extern int mbtowc(wchar_t *, const char *, size_t);
1920     extern char *mktemp(char *);
1921     extern long int mrand48(void);
1922     extern long int nrand48(unsigned short);
1923     extern char *ptsname(int);
1924     extern int putenv(char *);

```

```

1925     extern void qsort(void *, size_t, size_t, __compar_fn_t);
1926     extern int rand(void);
1927     extern int rand_r(unsigned int *);
1928     extern unsigned short *seed48(unsigned short);
1929     extern void srand48(long int);
1930     extern int unlockpt(int);
1931     extern size_t wcstombs(char *, const wchar_t *, size_t);
1932     extern int wctomb(char *, wchar_t);
1933     extern int system(const char *);
1934     extern void *calloc(size_t, size_t);
1935     extern void free(void *);
1936     extern char *initstate(unsigned int, char *, size_t);
1937     extern void *malloc(size_t);
1938     extern long int random(void);
1939     extern void *realloc(void *, size_t);
1940     extern char *setstate(char *);
1941     extern void srand(unsigned int);
1942     extern void srandrandom(unsigned int);
1943     extern double strtod(char *, char **);
1944     extern float strtof(const char *, char **);
1945     extern long int strtol(char *, char **, int);
1946     extern long double strtold(const char *, char **);
1947     extern long long int strtoll(const char *, char **, int);
1948     extern long long int strtoq(const char *, char **, int);
1949     extern unsigned long int strtoul(const char *, char **, int);
1950     extern unsigned long long int strtoull(const char *, char **, int);
1951     extern void _Exit(int);
1952     extern size_t __ctype_get_mb_cur_max(void);
1953     extern char **environ(void);
1955     extern char *realpath(const char *, char *);
1956     extern int setenv(const char *, const char *, int);
1957     extern int unsetenv(const char *);
1958     extern int getloadavg(double, int);
1959     extern int mkstemp64(char *);
1960     extern int posix_memalign(void **, size_t, size_t);
1961     extern int posix_openpt(int);

```

### 13.4.47 string.h

```

1962
1963     extern void *__mempcpy(void *, const void *, size_t);
1964     extern char *__stpncpy(char *, const char *);
1965     extern char *__strtok_r(char *, const char *, char **);
1966     extern void bcopy(void *, void *, size_t);
1967     extern void *memchr(void *, int, size_t);
1968     extern int memcmp(void *, void *, size_t);
1969     extern void *memcpy(void *, void *, size_t);
1970     extern void *memmem(const void *, size_t, const void *, size_t);
1971     extern void *memmove(void *, const void *, size_t);
1972     extern void *memset(void *, int, size_t);
1973     extern char *strcat(char *, const char *);
1974     extern char *strchr(char *, int);
1975     extern int strcmp(char *, char *);
1976     extern int strcoll(const char *, const char *);
1977     extern char *strcpy(char *, char *);
1978     extern size_t strcspn(const char *, const char *);
1979     extern char *strerror(int);
1980     extern size_t strlen(char *);
1981     extern char *strncat(char *, char *, size_t);
1982     extern int strncmp(char *, char *, size_t);
1983     extern char *strncpy(char *, char *, size_t);
1984     extern char *struprbrk(const char *, const char *);
1985     extern char *strrchr(char *, int);

```

```

1986 extern char *strsignal(int);
1987 extern size_t strspn(const char *, const char *);
1988 extern char *strstr(char *, char *);
1989 extern char *strtok(char *, const char *);
1990 extern size_t strxfrm(char *, const char *, size_t);
1991 extern int bcmp(void *, void *, size_t);
1992 extern void bzero(void *, size_t);
1993 extern int ffs(int);
1994 extern char *index(char *, int);
1995 extern void *memccpy(void *, const void *, int, size_t);
1996 extern char *rindex(char *, int);
1997 extern int strcasecmp(char *, char *);
1998 extern char *strdup(char *);
1999 extern int strncasecmp(char *, char *, size_t);
2000 extern char *strndup(const char *, size_t);
2001 extern size_t strnlen(const char *, size_t);
2002 extern char *strsep(char **, const char *);
2003 extern char *strerror_r(int, char *, size_t);
2004 extern char *strtok_r(char *, const char *, char **);
2005 extern char *strcasestr(const char *, const char *);
2006 extern char *stpcpy(char *, const char *);
2007 extern char *stpncpy(char *, const char *, size_t);
2008 extern void *memrchr(const void *, int, size_t);

```

### 13.4.48 sys/file.h

```

2009 #define LOCK_SH 1
2010 #define LOCK_EX 2
2011 #define LOCK_NB 4
2012 #define LOCK_UN 8
2013
2014
2015 extern int flock(int, int);

```

### 13.4.49 sys/ioctl.h

```

2016 struct winsize {
2017     unsigned short ws_row;
2018     unsigned short ws_col;
2019     unsigned short ws_xpixel;
2020     unsigned short ws_ypixel;
2021 };
2022
2023 extern int ioctl(int, unsigned long int, ...);

```

### 13.4.50 sys/ipc.h

```

2024 #define IPC_PRIVATE    ((key_t)0)
2025 #define IPC_RMID       0
2026 #define IPC_CREAT      00001000
2027 #define IPC_EXCL       00002000
2028 #define IPC_NOWAIT     00004000
2029 #define IPC_SET        1
2030 #define IPC_STAT       2
2031
2032
2033 extern key_t ftok(char *, int);

```

### 13.4.51 sys/mman.h

```

2034 #define MAP_FAILED      ((void*)-1)
2035 #define PROT_NONE       0x0
2036

```

```

2037      #define MAP_SHARED      0x01
2038      #define MAP_PRIVATE     0x02
2039      #define PROT_READ       0x1
2040      #define MAP_FIXED        0x10
2041      #define PROT_WRITE      0x2
2042      #define MAP_ANONYMOUS    0x20
2043      #define PROT_EXEC       0x4
2044      #define MS_ASYNC         1
2045      #define MS_INVALIDATE   2
2046      #define MS_SYNC          4
2047      #define MAP_ANON         MAP_ANONYMOUS
2048
2049      extern int msync(void *, size_t, int);
2050      extern int mlock(const void *, size_t);
2051      extern int mlockall(int);
2052      extern void *mmap(void *, size_t, int, int, int, off_t);
2053      extern int mprotect(void *, size_t, int);
2054      extern int munlock(const void *, size_t);
2055      extern int munlockall(void);
2056      extern int munmap(void *, size_t);
2057      extern void *mmap64(void *, size_t, int, int, int, off64_t);
2058      extern int shm_open(const char *, int, mode_t);
2059      extern int shm_unlink(const char *);

```

### 13.4.52 sys/msg.h

```

2060
2061      #define MSG_NOERROR      010000
2062
2063      extern int msgctl(int, int, struct msqid_ds *);
2064      extern int msgget(key_t, int);
2065      extern int msgrcv(int, void *, size_t, long int, int);
2066      extern int msgsnd(int, const void *, size_t, int);

```

### 13.4.53 sys/param.h

```

2067
2068      #define NOFILE    256
2069      #define MAXPATHLEN 4096

```

### 13.4.54 sys/poll.h

```

2070
2071      #define POLLIN   0x0001
2072      #define POLLPRI  0x0002
2073      #define POLLOUT  0x0004
2074      #define POLLERR  0x0008
2075      #define POLLHUP  0x0010
2076      #define POLLNVAL 0x0020
2077
2078      struct pollfd {
2079          int fd;
2080          short events;
2081          short revents;
2082      };
2083      typedef unsigned long int nfds_t;

```

### 13.4.55 sys/resource.h

```

2084
2085      #define RUSAGE_CHILDREN (-1)
2086      #define RUSAGE_BOTH      (-2)
2087      #define RLIM_INFINITY   (~0UL)

```

```

2088 #define RLIM_SAVED_CUR -1
2089 #define RLIM_SAVED_MAX -1
2090 #define RLIMIT_CPU 0
2091 #define RUSAGE_SELF 0
2092 #define RLIMIT_FSIZE 1
2093 #define RLIMIT_DATA 2
2094 #define RLIMIT_STACK 3
2095 #define RLIMIT_CORE 4
2096 #define RLIMIT_NOFILE 7
2097 #define RLIMIT_AS 9
2098
2099 typedef unsigned long int rlim_t;
2100 typedef unsigned long long int rlim64_t;
2101 typedef int __rlimit_resource_t;
2102
2103 struct rlimit {
2104     rlim_t rlim_cur;
2105     rlim_t rlim_max;
2106 };
2107 struct rlimit64 {
2108     rlim64_t rlim_cur;
2109     rlim64_t rlim_max;
2110 };
2111
2112 struct rusage {
2113     struct timeval ru_utime;
2114     struct timeval ru_stime;
2115     long int ru_maxrss;
2116     long int ru_ixrss;
2117     long int ru_idrss;
2118     long int ru_isrss;
2119     long int ru_minflt;
2120     long int ru_majflt;
2121     long int ru_nswap;
2122     long int ru_inblock;
2123     long int ru_oublock;
2124     long int ru_msgrnd;
2125     long int ru_msgrcv;
2126     long int ru_nsignals;
2127     long int ru_nvcsw;
2128     long int ru_nivcsw;
2129 };
2130
2131 enum __priority_which {
2132     PRIO_PGRP = 1, PRIO_PROCESS = 2
2133 };
2134
2135 #define PRIO_PGRP PRIO_PGRP
2136 #define PRIO_PROCESS PRIO_PROCESS
2137 #define PRIO_USER PRIO_USER
2138
2139 typedef enum __priority_which __priority_which_t;
2140 extern int getpriority(__priority_which_t, id_t);
2141 extern int getrlimit64(id_t, struct rlimit64 *);
2142 extern int setpriority(__priority_which_t, id_t, int);
2143 extern int setrlimit(__rlimit_resource_t, const struct rlimit *);
2144 extern int setrlimit64(__rlimit_resource_t, const struct rlimit64 *);
2145 extern int getrlimit(__rlimit_resource_t, struct rlimit *);
2146 extern int getrusage(int, struct rusage *);

```

### 13.4.56 sys/sem.h

```

2147
2148 #define SEM_UNDO 0x1000

```

```

2149     #define GETPID 11
2150     #define GETVAL 12
2151     #define GETALL 13
2152     #define GETNCNT 14
2153     #define GETZCNT 15
2154     #define SETVAL 16
2155     #define SETALL 17
2156
2157     struct sembuf {
2158         short sem_num;
2159         short sem_op;
2160         short sem_flg;
2161     };
2162     extern int semctl(int, int, int, ...);
2163     extern int semget(key_t, int, int);
2164     extern int semop(int, struct sembuf *, size_t);

```

### 13.4.57 sys/shm.h

```

2165
2166     #define SHM_RDONLY      010000
2167     #define SHM_W      0200
2168     #define SHM_RND 0200000
2169     #define SHM_R      0400
2170     #define SHM_REMAP      040000
2171     #define SHM_LOCK      11
2172     #define SHM_UNLOCK     12
2173
2174     extern int __getpagesize(void);
2175     extern void *shmat(int, const void *, int);
2176     extern int shmctl(int, int, struct shmid_ds *);
2177     extern int shmdt(const void *);
2178     extern int shmget(key_t, size_t, int);

```

### 13.4.58 sys/socket.h

```

2179
2180     #define CMSG_LEN(len)      (CMSG_ALIGN(sizeof(struct cmsghdr))+(len))
2181     #define SCM_RIGHTS        0x01
2182     #define SOL_SOCKET        1
2183     #define SOMAXCONN        128
2184     #define SOL_RAW           255
2185     #define CMSG_ALIGN(len) \
2186             (((len)+sizeof(size_t)-1)&(size_t)~(sizeof(size_t)-1))
2187     #define CMSG_DATA(cmsg) \
2188             ((unsigned char *) (cmsg) + CMSG_ALIGN(sizeof(struct cmsghdr)))
2189     #define CMSG_SPACE(len) \
2190             (CMSG_ALIGN(sizeof(struct cmsghdr))+CMSG_ALIGN(len))
2191     #define CMSG_FIRSTHDR(msg) \
2192             ((msg)->msg_controllen >= sizeof(struct cmsghdr) ? \
2193                 (struct cmsghdr *)(msg)->msg_control : \
2194                 (struct cmsghdr *)NULL)
2195     #define CMSG_NXTHDR(mhdr,cmsg) \
2196             (((cmsg) == NULL) ? CMSG_FIRSTHDR(mhdr) : \
2197                 (((u_char *) (cmsg) + CMSG_ALIGN((cmsg)->cmsg_len)) \
2198                  + CMSG_ALIGN(sizeof(struct cmsghdr)) > \
2199                  (u_char *) ((mhdr)->msg_control) + (mhdr)->msg_controllen) ?
2200                 \
2201                     (struct cmsghdr *)NULL : \
2202                     (struct cmsghdr *)((u_char *) (cmsg) + \
2203                         CMSG_ALIGN((cmsg)->cmsg_len))))
2204
2205     struct linger {
2206         int l_onoff;

```

```

2207         int l_linger;
2208     };
2209     struct cmsghdr {
2210         size_t cmsg_len;
2211         int cmsg_level;
2212         int cmsg_type;
2213     };
2214     struct iovec {
2215         void *iov_base;
2216         size_t iov_len;
2217     };
2218
2219     typedef unsigned short sa_family_t;
2220     typedef unsigned int socklen_t;
2221
2222     struct sockaddr {
2223         sa_family_t sa_family;
2224         char sa_data[14];
2225     };
2226     struct sockaddr_storage {
2227         sa_family_t ss_family;
2228         __ss_aligntype __ss_align;
2229         char __ss_padding[(128 - (2 * sizeof(__ss_aligntype)))];
2230     };
2231
2232     struct msghdr {
2233         void *msg_name;
2234         int msg_namelen;
2235         struct iovec *msg_iov;
2236         size_t msg iovlen;
2237         void *msg_control;
2238         size_t msg_controllen;
2239         unsigned int msg_flags;
2240     };
2241
2242     #define AF_UNSPEC          0
2243     #define AF_UNIX           1
2244     #define AF_INET6          10
2245     #define AF_INET           2
2246
2247     #define PF_INET           AF_INET
2248     #define PF_INET6          AF_INET6
2249     #define PF_UNIX           AF_UNIX
2250     #define PF_UNSPEC          AF_UNSPEC
2251
2252     #define SOCK_STREAM        1
2253     #define SOCK_PACKET        10
2254     #define SOCK_DGRAM         2
2255     #define SOCK_RAW           3
2256     #define SOCK_RDM           4
2257     #define SOCK_SEQPACKET     5
2258
2259     #define SO_DEBUG           1
2260     #define SO_OOBINLINE       10
2261     #define SO_NO_CHECK        11
2262     #define SO_PRIORITY        12
2263     #define SO_LINGER          13
2264     #define SO_REUSEADDR       2
2265     #define SO_TYPE            3
2266     #define SO_ACCEPTCONN      30
2267     #define SO_ERROR           4
2268     #define SO_DONTROUTE       5
2269     #define SO_BROADCAST       6
2270     #define SO_SNDBUF          7

```

```

2271     #define SO_RCVBUF      8
2272     #define SO_KEEPALIVE    9
2273
2274     #define SIOCGIFCONF     0x8912
2275     #define SIOCGIFFLAGS    0x8913
2276     #define SIOCGIFADDR     0x8915
2277     #define SIOCGIFNETMASK  0x891b
2278
2279     #define SHUT_RD 0
2280     #define SHUT_WR 1
2281     #define SHUT_RDWR 2
2282     #define MSG_DONTROUTE 4
2283
2284     #define MSG_WAITALL    0x100
2285     #define MSG_TRUNC      0x20
2286     #define MSG_EOR        0x80
2287     #define MSG_OOB        1
2288     #define MSG_PEEK       2
2289     #define MSG_CTRUNC     8
2290
2291     extern int bind(int, const struct sockaddr *, socklen_t);
2292     extern int getnameinfo(const struct sockaddr *, socklen_t, char *,
2293                           socklen_t, char *, socklen_t, unsigned int);
2294     extern int getsockname(int, struct sockaddr *, socklen_t *);
2295     extern int listen(int, int);
2296     extern int setsockopt(int, int, int, const void *, socklen_t);
2297     extern int accept(int, struct sockaddr *, socklen_t *);
2298     extern int connect(int, const struct sockaddr *, socklen_t);
2299     extern ssize_t recv(int, void *, size_t, int);
2300     extern ssize_t recvfrom(int, void *, size_t, int, struct sockaddr *,
2301                            socklen_t *);
2302     extern ssize_t recvmsg(int, struct msghdr *, int);
2303     extern ssize_t send(int, const void *, size_t, int);
2304     extern ssize_t sendmsg(int, const struct msghdr *, int);
2305     extern ssize_t sendto(int, const void *, size_t, int,
2306                           const struct sockaddr *, socklen_t);
2307     extern int getpeername(int, struct sockaddr *, socklen_t *);
2308     extern int getsockopt(int, int, int, void *, socklen_t *);
2309     extern int shutdown(int, int);
2310     extern int socket(int, int, int);
2311     extern int socketpair(int, int, int, int);
2312     extern int socketmark(int);

```

### 13.4.59 sys/stat.h

```

2313
2314     #define S_ISBLK(m)      (((m)&S_IFMT)==S_IFBLK)
2315     #define S_ISCHR(m)      (((m)&S_IFMT)==S_IFCHR)
2316     #define S_ISDIR(m)      (((m)&S_IFMT)==S_IFDIR)
2317     #define S_ISFIFO(m)     (((m)&S_IFMT)==S_IFIFO)
2318     #define S_ISLNK(m)      (((m)&S_IFMT)==S_IFLNK)
2319     #define S_ISREG(m)      (((m)&S_IFMT)==S_IFREG)
2320     #define S_ISSOCK(m)     (((m)&S_IFMT)==S_IFSOCK)
2321     #define S_TYPEISMQ(buf) ((buf)->st_mode - (buf)->st_mode)
2322     #define S_TYPEISSEM(buf) ((buf)->st_mode - (buf)->st_mode)
2323     #define S_TYPEISSSH(buf) ((buf)->st_mode - (buf)->st_mode)
2324     #define S_IRWXU (S_IREAD|S_IWRITE|S_IEXEC)
2325     #define S_IROTH (S_IRGRP>>3)
2326     #define S_IRGRP (S_IRUSR>>3)
2327     #define S_IRWXO (S_IRWXG>>3)
2328     #define S_IRWXG (S_IRWXU>>3)
2329     #define S_IWOTH (S_IWGRP>>3)
2330     #define S_IWGRP (S_IWUSR>>3)
2331     #define S_IXOTH (S_IXGRP>>3)

```

```

2332 #define S_IXGRP (S_IXUSR>>3)
2333 #define S_ISVTX 01000
2334 #define S_IXUSR 0x0040
2335 #define S_IWUSR 0x0080
2336 #define S_IRUSR 0x0100
2337 #define S_ISGID 0x0400
2338 #define S_ISUID 0x0800
2339 #define S_IFIFO 0x1000
2340 #define S_IFCHR 0x2000
2341 #define S_IFDIR 0x4000
2342 #define S_IFBLK 0x6000
2343 #define S_IFREG 0x8000
2344 #define S_IFLNK 0xa000
2345 #define S_IFSOCK 0xc000
2346 #define S_IFMT 0xf000
2347 #define st_atime st_atim.tv_sec
2348 #define st_ctime st_ctim.tv_sec
2349 #define st_mtime st_mtim.tv_sec
2350 #define S_IREAD S_IRUSR
2351 #define S_IWRITE S_IWUSR
2352 #define S_IEXEC S_IXUSR
2353
2354 extern int __fxstat(int, int, struct stat *);
2355 extern int __fxstat64(int, int, struct stat64 *);
2356 extern int __lxstat(int, char *, struct stat *);
2357 extern int __lxstat64(int, const char *, struct stat64 *);
2358 extern int __xmknod(int, const char *, mode_t, dev_t *);
2359 extern int __xstat(int, const char *, struct stat *);
2360 extern int __xstat64(int, const char *, struct stat64 *);
2361 extern int mkfifo(const char *, mode_t);
2362 extern int chmod(const char *, mode_t);
2363 extern int fchmod(int, mode_t);
2364 extern mode_t umask(mode_t);

```

### 13.4.60 sys/statvfs.h

```

2365
2366 extern int fstatvfs(int, struct statvfs *);
2367 extern int fstatvfs64(int, struct statvfs64 *);
2368 extern int statvfs(const char *, struct statvfs *);
2369 extern int statvfs64(const char *, struct statvfs64 *);

```

### 13.4.61 sys/time.h

```

2370
2371 #define ITIMER_REAL 0
2372 #define ITIMER_VIRTUAL 1
2373 #define ITIMER_PROF 2
2374
2375 struct timezone {
2376     int tz_minuteswest;
2377     int tz_dsttime;
2378 };
2379
2380 typedef int __itimer_which_t;
2381
2382 struct timespec {
2383     time_t tv_sec;
2384     long int tv_nsec;
2385 };
2386
2387 struct timeval {
2388     time_t tv_sec;
2389     suseconds_t tv_usec;

```

```

2390     } ;
2391
2392     struct itimerval {
2393         struct timeval it_interval;
2394         struct timeval it_value;
2395     } ;
2396     extern int getitimer(__itimer_which_t, struct itimerval *);
2397     extern int setitimer(__itimer_which_t, const struct itimerval *,
2398                          struct itimerval *);
2399     extern int adjtime(const struct timeval *, struct timeval *);
2400     extern int gettimeofday(struct timeval *, struct timezone *);
2401     extern int utimes(const char *, const struct timeval *);

```

### 13.4.62 sys/timeb.h

```

2402
2403     struct timeb {
2404         time_t time;
2405         unsigned short millitm;
2406         short timezone;
2407         short dstflag;
2408     } ;
2409     extern int ftime(struct timeb *) ;

```

### 13.4.63 sys/times.h

```

2410
2411     struct tms {
2412         clock_t tms_utime;
2413         clock_t tms_stime;
2414         clock_t tms_cutime;
2415         clock_t tms_cstime;
2416     } ;
2417     extern clock_t times(struct tms *) ;

```

### 13.4.64 sys/types.h

```

2418
2419     #define FALSE    0
2420     #define TRUE     1
2421     #define FD_SETSIZE      1024
2422     #define FD_ZERO(fdsetp) bzero(fdsetp, sizeof(*(fdsetp)))
2423     #define FD_ISSET(d, set) \
2424         ((set)->fds_bits[((d)/(8*sizeof(long)))]&(1<<((d)%(8*sizeof(long)))) )
2425     #define FD_CLR(d, set)   \
2426         ((set)->fds_bits[((d)/(8*sizeof(long)))]&=~(1<<((d)%(8*sizeof(long)))) )
2427     #define FD_SET(d, set)   \
2428         ((set)->fds_bits[((d)/(8*sizeof(long)))]|=(1<<((d)%(8*sizeof(long)))) )
2429
2430     typedef signed char int8_t;
2431     typedef short int16_t;
2432     typedef int int32_t;
2433     typedef unsigned char u_int8_t;
2434     typedef unsigned short u_int16_t;
2435     typedef unsigned int u_int32_t;
2436     typedef unsigned int uid_t;
2437     typedef int pid_t;

```

```

2444     typedef long int off_t;
2445     typedef int key_t;
2446     typedef long int suseconds_t;
2447     typedef unsigned int u_int;
2448     typedef struct {
2449         int __val[2];
2450     } fsid_t;
2451     typedef unsigned int useconds_t;
2452     typedef unsigned long int blksize_t;
2453     typedef long int fd_mask;
2454     typedef int timer_t;
2455     typedef int clockid_t;
2456
2457     typedef unsigned int id_t;
2458
2459     typedef unsigned long long int ino64_t;
2460     typedef long long int loff_t;
2461     typedef unsigned long int blkcnt_t;
2462     typedef unsigned long int fsblkcnt_t;
2463     typedef unsigned long int fsfilcnt_t;
2464     typedef unsigned long long int blkcnt64_t;
2465     typedef unsigned long long int fsblkcnt64_t;
2466     typedef unsigned long long int fsfilcnt64_t;
2467     typedef unsigned char u_char;
2468     typedef unsigned short u_short;
2469     typedef unsigned long int u_long;
2470
2471     typedef unsigned long int ino_t;
2472     typedef unsigned int gid_t;
2473     typedef unsigned long long int dev_t;
2474     typedef unsigned int mode_t;
2475     typedef unsigned long int nlink_t;
2476     typedef char *caddr_t;
2477
2478     typedef struct {
2479         unsigned long int fds_bits[__FDSET_LONGS];
2480     } fd_set;
2481
2482     typedef long int clock_t;
2483     typedef long int time_t;

```

### 13.4.65 sys/uio.h

```

2484
2485     extern ssize_t readv(int, const struct iovec *, int);
2486     extern ssize_t writev(int, const struct iovec *, int);

```

### 13.4.66 sys/un.h

```

2487
2488     #define UNIX_PATH_MAX    108
2489
2490     struct sockaddr_un {
2491         sa_family_t sun_family;
2492         char sun_path[UNIX_PATH_MAX];
2493     };

```

### 13.4.67 sys/utsname.h

```

2494
2495     #define SYS_NMLN          65
2496
2497     struct utsname {

```

```

2498         char sysname[65];
2499         char nodename[65];
2500         char release[65];
2501         char version[65];
2502         char machine[65];
2503         char domainname[65];
2504     };
2505     extern int uname(struct utsname *);

```

### 13.4.68 sys/wait.h

```

2506 #define WIFSIGNALED(status)      (!WIFSTOPPED(status)
2507 && !WIFEXITED(status))
2508 #define WIFSTOPPED(status)      (((status) & 0xff) == 0x7f)
2509 #define WEXITSTATUS(status)    (((status) & 0xff00) >> 8)
2510 #define WTERMSIG(status)       ((status) & 0x7f)
2511 #define WCOREDUMP(status)      ((status) & 0x80)
2512 #define WIFEXITED(status)      (WTERMSIG(status) == 0)
2513 #define WNOHANG 0x00000001
2514 #define WUNTRACED 0x00000002
2515 #define WCOREFLAG 0x80
2516 #define WSTOPSIG(status)      WEXITSTATUS(status)
2517
2518 typedef enum {
2519     P_ALL, P_PID, P_PGID
2520 } idtype_t;
2521 extern pid_t wait(int *);
2522 extern pid_t waitpid(pid_t, int *, int);
2523 extern pid_t wait4(pid_t, int *, int, struct rusage *);
2524

```

### 13.4.69 syslog.h

```

2525
2526 #define LOG_EMERG      0
2527 #define LOG_PRIMASK    0x07
2528 #define LOG_ALERT      1
2529 #define LOG_CRIT       2
2530 #define LOG_ERR        3
2531 #define LOG_WARNING    4
2532 #define LOG_NOTICE     5
2533 #define LOG_INFO       6
2534 #define LOG_DEBUG      7
2535
2536 #define LOG_KERN       (0<<3)
2537 #define LOG_AUTHPRIV   (10<<3)
2538 #define LOG_FTP        (11<<3)
2539 #define LOG_USER       (1<<3)
2540 #define LOG_MAIL       (2<<3)
2541 #define LOG_DAEMON    (3<<3)
2542 #define LOG_AUTH      (4<<3)
2543 #define LOG_SYSLOG    (5<<3)
2544 #define LOG_LPR        (6<<3)
2545 #define LOG_NEWS       (7<<3)
2546 #define LOG_UUCP       (8<<3)
2547 #define LOG_CRON       (9<<3)
2548 #define LOG_FACMASK   0x03f8
2549
2550 #define LOG_LOCAL0     (16<<3)
2551 #define LOG_LOCAL1     (17<<3)
2552 #define LOG_LOCAL2     (18<<3)
2553 #define LOG_LOCAL3     (19<<3)
2554 #define LOG_LOCAL4     (20<<3)
2555 #define LOG_LOCAL5     (21<<3)

```

```

2556 #define LOG_LOCAL6      (22<<3)
2557 #define LOG_LOCAL7      (23<<3)
2558
2559 #define LOG_UPTO(pri)   (((1 << ((pri)+1)) - 1)
2560 #define LOG_MASK(pri)   (1 << (pri))
2561
2562 #define LOG_PID 0x01
2563 #define LOG_CONS 0x02
2564 #define LOG_ODELAY 0x04
2565 #define LOG_NDELAY 0x08
2566 #define LOG_NOWAIT 0x10
2567 #define LOG_PERROR 0x20
2568
2569 extern void closelog(void);
2570 extern void openlog(const char *, int, int);
2571 extern int setlogmask(int);
2572 extern void syslog(int, const char *, ...);
2573 extern void vsyslog(int, const char *, va_list);

```

### 13.4.70 termios.h

```

2574
2575 #define TCIFLUSH      0
2576 #define TCOOFF       0
2577 #define TCSANOW      0
2578 #define BS0          0000000
2579 #define CR0          0000000
2580 #define FF0          0000000
2581 #define NL0          0000000
2582 #define TAB0         0000000
2583 #define VT0          0000000
2584 #define OPOST         0000001
2585 #define OCRNL        0000010
2586 #define ONOCR        0000020
2587 #define ONLRET        0000040
2588 #define OFILL         0000100
2589 #define OFDEL         0000200
2590 #define NL1          0000400
2591 #define TCOFLUSH      1
2592 #define TCOON         1
2593 #define TCSADRAIN     1
2594 #define TCIOFF        2
2595 #define TCIOFLUSH     2
2596 #define TCSAFLUSH     2
2597 #define TCION         3
2598
2599 typedef unsigned int speed_t;
2600 typedef unsigned char cc_t;
2601 typedef unsigned int tcflag_t;
2602
2603 #define NCCS          32
2604
2605 struct termios {
2606     tcflag_t c_iflag;
2607     tcflag_t c_oflag;
2608     tcflag_t c_cflag;
2609     tcflag_t c_lflag;
2610     cc_t c_line;
2611     cc_t c_cc[NCCS];
2612     speed_t c_ispeed;
2613     speed_t c_ospeed;
2614 };
2615
2616 #define VINTR        0

```

```

2617      #define VQUIT    1
2618      #define VLNEXTT 15
2619      #define VERASE   2
2620      #define VKILL    3
2621      #define VEOF     4
2622
2623      #define IGNBRK   0000001
2624      #define BRKINT   0000002
2625      #define IGNPAR   0000004
2626      #define PARMRK   0000010
2627      #define INPCK    0000020
2628      #define ISTRIP   0000040
2629      #define INLCR    0000100
2630      #define IGNCR    0000200
2631      #define ICRNL    0000400
2632      #define IXANY    0004000
2633      #define IMAXBEL  0020000
2634
2635      #define CS5     0000000
2636
2637      #define ECHO    0000010
2638
2639      #define B0      0000000
2640      #define B50     0000001
2641      #define B75      0000002
2642      #define B110     0000003
2643      #define B134     0000004
2644      #define B150     0000005
2645      #define B200     0000006
2646      #define B300     0000007
2647      #define B600     0000010
2648      #define B1200    0000011
2649      #define B1800    0000012
2650      #define B2400    0000013
2651      #define B4800    0000014
2652      #define B9600    0000015
2653      #define B19200   0000016
2654      #define B38400   0000017
2655
2656      extern speed_t cfgetispeed(const struct termios *);
2657      extern speed_t cfgetospeed(const struct termios *);
2658      extern void cfmakeraw(struct termios *);
2659      extern int cfsetispeed(struct termios *, speed_t);
2660      extern int cfsetospeed(struct termios *, speed_t);
2661      extern int cfsetspeed(struct termios *, speed_t);
2662      extern int tcflow(int, int);
2663      extern int tcflush(int, int);
2664      extern pid_t tcgetsid(int);
2665      extern int tcsendbreak(int, int);
2666      extern int tcsetattr(int, int, const struct termios *);
2667      extern int tcdrain(int);
2668      extern int tcgetattr(int, struct termios *);

```

### 13.4.71 time.h

```

2669
2670      #define CLK_TCK ((clock_t) __sysconf(2))
2671      #define CLOCK_REALTIME 0
2672      #define TIMER_ABSTIME 1
2673      #define CLOCKS_PER_SEC 10000001
2674
2675      struct tm {
2676          int tm_sec;
2677          int tm_min;

```

```

2678     int tm_hour;
2679     int tm_mday;
2680     int tm_mon;
2681     int tm_year;
2682     int tm_wday;
2683     int tm_yday;
2684     int tm_isdst;
2685     long int tm_gmtoff;
2686     char *tm_zone;
2687 };
2688 struct itimerspec {
2689     struct timespec it_interval;
2690     struct timespec it_value;
2691 };
2692
2693 extern int __daylight(void);
2694 extern long int __timezone(void);
2695 extern char *__tzname(void);
2696 extern char *asctime(const struct tm *);
2697 extern clock_t clock(void);
2698 extern char *ctime(const time_t *);
2699 extern char *ctime_r(const time_t *, char *);
2700 extern double difftime(time_t, time_t);
2701 extern struct tm *getdate(const char *);
2702 extern int getdate_err(void);
2703 extern struct tm *gmtime(const time_t *);
2704 extern struct tm *localtime(const time_t *);
2705 extern time_t mktime(struct tm *);
2706 extern int stime(const time_t *);
2707 extern size_t strftime(char *, size_t, const char *, const struct tm *);
2708 extern char *strptime(const char *, const char *, struct tm *);
2709 extern time_t time(time_t *);
2710 extern int nanosleep(const struct timespec *, struct timespec *);
2711 extern int daylight(void);
2712 extern long int timezone(void);
2713 extern char *tzname(void);
2714 extern void tzset(void);
2715 extern char *asctime_r(const struct tm *, char *);
2716 extern struct tm *gmtime_r(const time_t *, struct tm *);
2717 extern struct tm *localtime_r(const time_t *, struct tm *);
2718 extern int clock_getcpuclockid(pid_t, clockid_t *);
2719 extern int clock_getres(clockid_t, struct timespec *);
2720 extern int clock_gettime(clockid_t, struct timespec *);
2721 extern int clock_nanosleep(clockid_t, int, const struct timespec *,
2722                             struct timespec *);
2723 extern int clock_settime(clockid_t, const struct timespec *);
2724 extern int timer_create(clockid_t, struct sigevent *, timer_t *);
2725 extern int timer_delete(timer_t);
2726 extern int timer_getoverrun(timer_t);
2727 extern int timer_gettime(timer_t, struct itimerspec *);
2728 extern int timer_settime(timer_t, int, const struct itimerspec *,
2729                         struct itimerspec *);

```

### 13.4.72 ucontext.h

```

2730
2731 extern int getcontext(ucontext_t *);
2732 extern int makecontext(ucontext_t *, void (*func) (void)
2733                           , int, ...);
2734 extern int setcontext(const struct ucontext *);
2735 extern int swapcontext(ucontext_t *, const struct ucontext *);

```

### 13.4.73 ulimit.h

```

2736
2737     #define UL_GETFSIZE      1
2738     #define UL_SETFSIZE      2
2739
2740     extern long int ulimit(int, ...);

```

### 13.4.74 unistd.h

```

2741
2742     #define SEEK_SET          0
2743     #define STDIN_FILENO       0
2744     #define SEEK_CUR           1
2745     #define STDOUT_FILENO      1
2746     #define SEEK_END           2
2747     #define STDERR_FILENO      2
2748
2749     typedef long long int off64_t;
2750
2751     #define F_OK        0
2752     #define X_OK        1
2753     #define W_OK        2
2754     #define R_OK        4
2755
2756     #define _POSIX_VDISABLE '\0'
2757     #define _POSIX_CHOWN_RESTRICTED 1
2758     #define _POSIX_JOB_CONTROL    1
2759     #define _POSIX_NO_TRUNC       1
2760     #define _POSIX_SHELL          1
2761     #define _POSIX_FSYNC          200112
2762     #define _POSIX_MAPPED_FILES   200112
2763     #define _POSIX_MEMLOCK         200112
2764     #define _POSIX_MEMLOCK_RANGE   200112
2765     #define _POSIX_MEMORY_PROTECTION 200112
2766     #define _POSIX_SEMAPHORES      200112
2767     #define _POSIX_SHARED_MEMORY_OBJECTS 200112
2768     #define _POSIX_TIMERS         200112
2769     #define _POSIX2_C_BIND        200112L
2770     #define _POSIX_THREADS         200112L
2771
2772     #define _PC_LINK_MAX          0
2773     #define _PC_MAX_CANON        1
2774     #define _PC_ASYNC_IO          10
2775     #define _PC_PRIO_IO          11
2776     #define _PC_FILESIZEBITS      13
2777     #define _PC_REC_INCR_XFER_SIZE 14
2778     #define _PC_REC_MIN_XFER_SIZE 16
2779     #define _PC_REC_XFER_ALIGN    17
2780     #define _PC_ALLOC_SIZE_MIN    18
2781     #define _PC_MAX_INPUT         2
2782     #define _PC_2_SYMLINKS        20
2783     #define _PC_NAME_MAX          3
2784     #define _PC_PATH_MAX          4
2785     #define _PC_PIPE_BUF           5
2786     #define _PC_CHOWN_RESTRICTED  6
2787     #define _PC_NO_TRUNC          7
2788     #define _PC_VDISABLE          8
2789     #define _PC_SYNC_IO           9
2790
2791     #define _SC_ARG_MAX           0
2792     #define _SC_CHILD_MAX          1
2793     #define _SC_PRIORITY_SCHEDULING 10
2794     #define _SC_TIMERS            11

```

```

2795 #define _SC_ASYNCHRONOUS_IO      12
2796 #define _SC_XBS5_ILP32_OFF32    125
2797 #define _SC_XBS5_ILP32_OFFBIG   126
2798 #define _SC_XBS5_LP64_OFF64    127
2799 #define _SC_XBS5_LP64_OFFBIG   128
2800 #define _SC_XOPEN_LEGACY       129
2801 #define _SC_PRIORITIZED_IO     13
2802 #define _SC_XOPEN_REALTIME     130
2803 #define _SC_XOPEN_REALTIME_THREADS 131
2804 #define _SC_ADVISORY_INFO      132
2805 #define _SC_BARRIERS          133
2806 #define _SC_CLOCK_SELECTION    137
2807 #define _SC_CPUTIME            138
2808 #define _SC_THREAD_CPUTIME     139
2809 #define _SC_SYNCHRONIZED_IO    14
2810 #define _SC_MONOTONIC_CLOCK    149
2811 #define _SC_FSYNC              15
2812 #define _SC_READER_WRITER_LOCKS 153
2813 #define _SC_SPIN_LOCKS         154
2814 #define _SC_REGEXP              155
2815 #define _SC_SHELL               157
2816 #define _SC_SPAWN              159
2817 #define _SC_MAPPED_FILES        16
2818 #define _SC_SPORADIC_SERVER    160
2819 #define _SC_THREAD_SPORADIC_SERVER 161
2820 #define _SC_TIMEOUTS            164
2821 #define _SC_TYPED_MEMORY_OBJECTS 165
2822 #define _SC_2_PBS_ACCOUNTING    169
2823 #define _SC_MEMLOCK              17
2824 #define _SC_2_PBS_LOCATE        170
2825 #define _SC_2_PBS_MESSAGE       171
2826 #define _SC_2_PBS_TRACK         172
2827 #define _SC_SYMLOOP_MAX         173
2828 #define _SC_2_PBS_CHECKPOINT    175
2829 #define _SC_V6_ILP32_OFF32      176
2830 #define _SC_V6_ILP32_OFFBIG     177
2831 #define _SC_V6_LP64_OFF64       178
2832 #define _SC_V6_LP64_OFFBIG      179
2833 #define _SC_MEMLOCK_RANGE        18
2834 #define _SC_HOST_NAME_MAX       180
2835 #define _SC_TRACE                181
2836 #define _SC_TRACE_EVENT_FILTER   182
2837 #define _SC_TRACE_INHERIT        183
2838 #define _SC_TRACE_LOG             184
2839 #define _SC_MEMORY_PROTECTION    19
2840 #define _SC_CLK_TCK                2
2841 #define _SC_MESSAGE_PASSING      20
2842 #define _SC_SEMAPHORES            21
2843 #define _SC_SHARED_MEMORY_OBJECTS 22
2844 #define _SC_AIO_LISTIO_MAX        23
2845 #define _SC_AIO_MAX                24
2846 #define _SC_AIO_PRIO_DELTA_MAX   25
2847 #define _SC_DELAYTIMER_MAX        26
2848 #define _SC_MQ_OPEN_MAX           27
2849 #define _SC_MQ_PRIO_MAX           28
2850 #define _SC_VERSION                29
2851 #define _SC_NGROUPS_MAX            3
2852 #define _SC_PAGESIZE               30
2853 #define _SC_PAGE_SIZE               30
2854 #define _SC_RTSIG_MAX                31
2855 #define _SC_SEM_NSEMS_MAX           32
2856 #define _SC_SEM_VALUE_MAX           33
2857 #define _SC_SIGQUEUE_MAX            34
2858 #define _SC_TIMER_MAX                35

```

```

2859      #define _SC_BC_BASE_MAX 36
2860      #define _SC_BC_DIM_MAX 37
2861      #define _SC_BC_SCALE_MAX 38
2862      #define _SC_BC_STRING_MAX 39
2863      #define _SC_OPEN_MAX 4
2864      #define _SC_COLL_WEIGHTS_MAX 40
2865      #define _SC_EXPR_NEST_MAX 42
2866      #define _SC_LINE_MAX 43
2867      #define _SC_RE_DUP_MAX 44
2868      #define _SC_2_VERSION 46
2869      #define _SC_2_C_BIND 47
2870      #define _SC_2_C_DEV 48
2871      #define _SC_2_FORT_DEV 49
2872      #define _SC_STREAM_MAX 5
2873      #define _SC_2_FORT_RUN 50
2874      #define _SC_2_SW_DEV 51
2875      #define _SC_2_LOCALEDEF 52
2876      #define _SC_TZNAME_MAX 6
2877      #define _SC_IOV_MAX 60
2878      #define _SC_THREADS 67
2879      #define _SC_THREAD_SAFE_FUNCTIONS 68
2880      #define _SC_GETGR_R_SIZE_MAX 69
2881      #define _SC_JOB_CONTROL 7
2882      #define _SC_GETPW_R_SIZE_MAX 70
2883      #define _SC_LOGIN_NAME_MAX 71
2884      #define _SC_TTY_NAME_MAX 72
2885      #define _SC_THREAD_DESTRUCTOR_ITERATIONS 73
2886      #define _SC_THREAD_KEYS_MAX 74
2887      #define _SC_THREAD_STACK_MIN 75
2888      #define _SC_THREAD_THREADS_MAX 76
2889      #define _SC_THREAD_ATTR_STACKADDR 77
2890      #define _SC_THREAD_ATTR_STACKSIZE 78
2891      #define _SC_THREAD_PRIORITY_SCHEDULING 79
2892      #define _SC_SAVED_IDS 8
2893      #define _SC_THREAD_PRIO_INHERIT 80
2894      #define _SC_THREAD_PRIO_PROTECT 81
2895      #define _SC_THREAD_PROCESS_SHARED 82
2896      #define _SC_ATEXIT_MAX 87
2897      #define _SC_PASS_MAX 88
2898      #define _SC_XOPEN_VERSION 89
2899      #define _SC_REALTIME_SIGNALS 9
2900      #define _SC_XOPEN_UNIX 91
2901      #define _SC_XOPEN_CRYPT 92
2902      #define _SC_XOPEN_ENH_I18N 93
2903      #define _SC_XOPEN_SHM 94
2904      #define _SC_2_CHAR_TERM 95
2905      #define _SC_2_C_VERSION 96
2906      #define _SC_2_UPE 97
2907
2908      #define _CS_PATH 0
2909      #define _POSIX_REGEX 1
2910      #define _CS_XBS5_ILP32_OFF32_CFLAGS 1100
2911      #define _CS_XBS5_ILP32_OFF32_LDFLAGS 1101
2912      #define _CS_XBS5_ILP32_OFF32_LIBS 1102
2913      #define _CS_XBS5_ILP32_OFF32_LINTFLAGS 1103
2914      #define _CS_XBS5_ILP32_OFFBIG_CFLAGS 1104
2915      #define _CS_XBS5_ILP32_OFFBIG_LDFLAGS 1105
2916      #define _CS_XBS5_ILP32_OFFBIG_LIBS 1106
2917      #define _CS_XBS5_ILP32_OFFBIG_LINTFLAGS 1107
2918      #define _CS_XBS5_LP64_OFF64_CFLAGS 1108
2919      #define _CS_XBS5_LP64_OFF64_LDFLAGS 1109
2920      #define _CS_XBS5_LP64_OFF64_LIBS 1110
2921      #define _CS_XBS5_LP64_OFF64_LINTFLAGS 1111
2922      #define _CS_XBS5_LP64_CFLAGS 1112

```

```

2923 #define _CS_XBS5_LPBIK_OFFBIG_LDFLAGS    1113
2924 #define _CS_XBS5_LPBIK_OFFBIG_LIBS      1114
2925 #define _CS_XBS5_LPBIK_OFFBIG_LINTFLAGS 1115
2926
2927 #define _XOPEN_XPG4      1
2928
2929 #define F_UNLOCK 0
2930 #define F_LOCK   1
2931 #define F_TLOCK  2
2932 #define F_TEST   3
2933
2934 extern char **__environ(void);
2935 extern pid_t __getpgid(pid_t);
2936 extern void _exit(int);
2937 extern int acct(const char *);
2938 extern unsigned int alarm(unsigned int);
2939 extern int chown(const char *, uid_t, gid_t);
2940 extern int chroot(const char *);
2941 extern size_t confstr(int, char *, size_t);
2942 extern int creat(const char *, mode_t);
2943 extern int creat64(const char *, mode_t);
2944 extern char *ctermid(char *);
2945 extern char *cuserid(char *);
2946 extern int daemon(int, int);
2947 extern int execl(const char *, const char *, ...);
2948 extern int execle(const char *, const char *, ...);
2949 extern int execlp(const char *, const char *, ...);
2950 extern int execv(const char *, char *const);
2951 extern int execvp(const char *, char *const);
2952 extern int fdatasync(int);
2953 extern int ftruncate64(int, off64_t);
2954 extern long int gethostid(void);
2955 extern char *getlogin(void);
2956 extern int getlogin_r(char *, size_t);
2957 extern int getopt(int, char *const, const char *);
2958 extern pid_t getpgrp(void);
2959 extern pid_t getsid(pid_t);
2960 extern char *getwd(char *);
2961 extern int lockf(int, int, off_t);
2962 extern int mkstemp(char *);
2963 extern int nice(int);
2964 extern char *optarg(void);
2965 extern int opterr(void);
2966 extern int optind(void);
2967 extern int optopt(void);
2968 extern int rename(const char *, const char *);
2969 extern int setegid(gid_t);
2970 extern int seteuid(uid_t);
2971 extern int sethostname(const char *, size_t);
2972 extern int setpgrp(void);
2973 extern void swab(const void *, void *, ssize_t);
2974 extern void sync(void);
2975 extern pid_t tcgetpgrp(int);
2976 extern int tcsetpgrp(int, pid_t);
2977 extern int truncate(const char *, off_t);
2978 extern int truncate64(const char *, off64_t);
2979 extern char *ttyname(int);
2980 extern unsigned int ualarm(useconds_t, useconds_t);
2981 extern int usleep(useconds_t);
2982 extern int close(int);
2983 extern int fsync(int);
2984 extern off_t lseek(int, off_t, int);
2985 extern int open(const char *, int, ...);
2986 extern int pause(void);

```

```

2987     extern ssize_t read(int, void *, size_t);
2988     extern ssize_t write(int, const void *, size_t);
2989     extern char *crypt(char *, char *);
2990     extern void encrypt(char *, int);
2991     extern void setkey(const char *);
2992     extern int access(const char *, int);
2993     extern int brk(void *);
2994     extern int chdir(const char *);
2995     extern int dup(int);
2996     extern int dup2(int, int);
2997     extern int execve(const char *, char *const, char *const);
2998     extern int fchdir(int);
2999     extern int fchown(int, uid_t, gid_t);
3000     extern pid_t fork(void);
3001     extern gid_t getegid(void);
3002     extern uid_t geteuid(void);
3003     extern gid_t getgid(void);
3004     extern int getgroups(int, gid_t);
3005     extern int gethostname(char *, size_t);
3006     extern pid_t getpgid(pid_t);
3007     extern pid_t getpid(void);
3008     extern uid_t getuid(void);
3009     extern int lchown(const char *, uid_t, gid_t);
3010     extern int link(const char *, const char *);
3011     extern int mkdir(const char *, mode_t);
3012     extern long int pathconf(const char *, int);
3013     extern int pipe(int);
3014     extern int readlink(const char *, char *, size_t);
3015     extern int rmdir(const char *);
3016     extern void *sbrk(ptrdiff_t);
3017     extern int select(int, fd_set *, fd_set *, fd_set *, struct timeval *);
3018     extern int setgid(gid_t);
3019     extern int setpgid(pid_t, pid_t);
3020     extern int setregid(gid_t, gid_t);
3021     extern int setreuid(uid_t, uid_t);
3022     extern pid_t setsid(void);
3023     extern int setuid(uid_t);
3024     extern unsigned int sleep(unsigned int);
3025     extern int symlink(const char *, const char *);
3026     extern long int sysconf(int);
3027     extern int unlink(const char *);
3028     extern pid_t vfork(void);
3029     extern ssize_t pread(int, void *, size_t, off_t);
3030     extern ssize_t pwrite(int, const void *, size_t, off_t);
3031     extern char **_environ(void);
3032     extern long int fpathconf(int, int);
3033     extern int ftruncate(int, off_t);
3034     extern char *getcwd(char *, size_t);
3035     extern int getpagesize(void);
3036     extern pid_t getppid(void);
3037     extern int isatty(int);
3038     extern loff_t lseek64(int, loff_t, int);
3039     extern int open64(const char *, int, ...);
3040     extern ssize_t pread64(int, void *, size_t, off64_t);
3041     extern ssize_t pwrite64(int, const void *, size_t, off64_t);
3042     extern int ttynname_r(int, char *, size_t);

```

### 13.4.75 utime.h

```

3043
3044     struct utimbuf {
3045         time_t actime;
3046         time_t modtime;
3047     };

```

```
3048 extern int utime(const char *, const struct utimbuf *);
```

### 13.4.76 utmp.h

```
3049
3050     #define UT_HOSTSIZE      256
3051     #define UT_LINESIZE       32
3052     #define UT_NAMESIZE       32
3053
3054     struct exit_status {
3055         short e_termination;
3056         short e_exit;
3057     };
3058
3059     #define EMPTY      0
3060     #define RUN_LVL    1
3061     #define BOOT_TIME   2
3062     #define NEW_TIME    3
3063     #define OLD_TIME    4
3064     #define INIT_PROCESS 5
3065     #define LOGIN_PROCESS 6
3066     #define USER_PROCESS  7
3067     #define DEAD_PROCESS  8
3068     #define ACCOUNTING  9
3069
3070     extern void endutent(void);
3071     extern struct utmp *getutent(void);
3072     extern void setutent(void);
3073     extern int getutent_r(struct utmp *, struct utmp **);
3074     extern int utmpname(const char *);
3075     extern int login_tty(int);
3076     extern void login(const struct utmp *);
3077     extern int logout(const char *);
3078     extern void logwtmp(const char *, const char *, const char *);
```

### 13.4.77 utmpx.h

```
3079
3080     extern void endutxent(void);
3081     extern struct utmpx *getutxent(void);
3082     extern struct utmpx *getutxid(const struct utmpx *);
3083     extern struct utmpx *getutxline(const struct utmpx *);
3084     extern struct utmpx *pututxline(const struct utmpx *);
3085     extern void setutxent(void);
```

### 13.4.78 wchar.h

```
3086
3087     #define WEOF      (0xfffffffffu)
3088     #define WCHAR_MAX    0x7FFFFFFF
3089     #define WCHAR_MIN    0x80000000
3090
3091     extern double __wcstod_internal(const wchar_t *, wchar_t **, int);
3092     extern float __wcstof_internal(const wchar_t *, wchar_t **, int);
3093     extern long int __wcstol_internal(const wchar_t *, wchar_t **, int,
3094                                         int);
3095     extern long double __wcstold_internal(const wchar_t *, wchar_t **, int);
3096     extern unsigned long int __wcstoul_internal(const wchar_t *, wchar_t *
3097                                         *,
3098                                         int, int);
3099     extern wchar_t *wcscat(wchar_t *, const wchar_t *);
3100     extern wchar_t *wcschr(const wchar_t *, wchar_t);
3101     extern int wcscmp(const wchar_t *, const wchar_t *);
```

```

3102     extern int wcscoll(const wchar_t *, const wchar_t *);
3103     extern wchar_t *wcscpy(wchar_t *, const wchar_t *);
3104     extern size_t wcscspn(const wchar_t *, const wchar_t *);
3105     extern wchar_t *wcsdup(const wchar_t *);
3106     extern wchar_t *wcsncat(wchar_t *, const wchar_t *, size_t);
3107     extern int wcsncmp(const wchar_t *, const wchar_t *, size_t);
3108     extern wchar_t *wcsncpy(wchar_t *, const wchar_t *, size_t);
3109     extern wchar_t *wcspbrk(const wchar_t *, const wchar_t *);
3110     extern wchar_t *wcsrchr(const wchar_t *, wchar_t);
3111     extern size_t wcspn(const wchar_t *, const wchar_t *);
3112     extern wchar_t *wcssstr(const wchar_t *, const wchar_t *);
3113     extern wchar_t *wcstok(wchar_t *, const wchar_t *, wchar_t **);
3114     extern int wcswidth(const wchar_t *, size_t);
3115     extern size_t wcsxfrm(wchar_t *, const wchar_t *, size_t);
3116     extern int wctob(wint_t);
3117     extern int wcwidth(wchar_t);
3118     extern wchar_t *wmemchr(const wchar_t *, wchar_t, size_t);
3119     extern int wmemcmp(const wchar_t *, const wchar_t *, size_t);
3120     extern wchar_t *wmemcpy(wchar_t *, const wchar_t *, size_t);
3121     extern wchar_t *wmemmove(wchar_t *, const wchar_t *, size_t);
3122     extern wchar_t *wmemset(wchar_t *, wchar_t, size_t);
3123     extern size_t mbrlen(const char *, size_t, mbstate_t *);
3124     extern size_t mbrtowc(wchar_t *, const char *, size_t, mbstate_t *);
3125     extern int mbsinit(const mbstate_t *);
3126     extern size_t mbsnrtowcs(wchar_t *, const char **, size_t, size_t,
3127                             mbstate_t *);
3128     extern size_t mbsrtowcs(wchar_t *, const char **, size_t, mbstate_t *);
3129     extern wchar_t *wcpcpy(wchar_t *, const wchar_t *);
3130     extern wchar_t *wcpncpy(wchar_t *, const wchar_t *, size_t);
3131     extern size_t wcrtomb(char *, wchar_t, mbstate_t *);
3132     extern size_t wcslen(const wchar_t *);
3133     extern size_t wcsnrtombs(char *, const wchar_t **, size_t, size_t,
3134                             mbstate_t *);
3135     extern size_t wcsrtnombs(char *, const wchar_t **, size_t, mbstate_t *);
3136     extern double wcstod(const wchar_t *, wchar_t **);
3137     extern float wcstof(const wchar_t *, wchar_t **);
3138     extern long int wcstol(const wchar_t *, wchar_t **, int);
3139     extern long double wcstold(const wchar_t *, wchar_t **);
3140     extern long long int wcstoq(const wchar_t *, wchar_t **, int);
3141     extern unsigned long int wcstoul(const wchar_t *, wchar_t **, int);
3142     extern unsigned long long int wcstouq(const wchar_t *, wchar_t **, int);
3143     extern wchar_t *wcswcs(const wchar_t *, const wchar_t *);
3144     extern int wcscasecmp(const wchar_t *, const wchar_t *);
3145     extern int wcsncasecmp(const wchar_t *, const wchar_t *, size_t);
3146     extern size_t wcsnlen(const wchar_t *, size_t);
3147     extern long long int wcstoll(const wchar_t *, wchar_t **, int);
3148     extern unsigned long long int wcstoull(const wchar_t *, wchar_t **, int);
3149     extern wint_t btowc(int);
3150     extern wint_t fgetwc(FILE *);
3151     extern wint_t fgetwc_unlocked(FILE *);
3152     extern wchar_t *fgetws(wchar_t *, int, FILE *);
3153     extern wint_t fputwc(wchar_t, FILE *);
3154     extern int fputws(const wchar_t *, FILE *);
3155     extern int fwide(FILE *, int);
3156     extern int fwprintf(FILE *, const wchar_t *, ...);
3157     extern int fwscanf(FILE *, const wchar_t *, ...);
3158     extern wint_t getwc(FILE *);
3159     extern wint_t getwchar(void);
3160     extern wint_t putwc(wchar_t, FILE *);
3161     extern wint_t putwchar(wchar_t);
3162     extern int swprintf(wchar_t *, size_t, const wchar_t *, ...);
3163     extern int swscanf(const wchar_t *, const wchar_t *, ...);
3164     extern wint_t ungetwc(wint_t, FILE *);
3165     extern int vfwprintf(FILE *, const wchar_t *, va_list);

```

```

3166 extern int vfwscanf(FILE *, const wchar_t *, va_list);
3167 extern int vswprintf(wchar_t *, size_t, const wchar_t *, va_list);
3168 extern int vsscanf(const wchar_t *, const wchar_t *, va_list);
3169 extern int vwprintf(const wchar_t *, va_list);
3170 extern int vwscanf(const wchar_t *, va_list);
3171 extern size_t wcsftime(wchar_t *, size_t, const wchar_t *,
3172                         const struct tm *);
3173 extern int wprintf(const wchar_t *, ...);
3174 extern int wscanf(const wchar_t *, ...);

```

### 13.4.79 wctype.h

```

3175
3176     typedef unsigned long int wctype_t;
3177     typedef unsigned int wint_t;
3178     typedef const int32_t *wctrans_t;
3179     typedef struct {
3180         int count;
3181         wint_t value;
3182     } __mbstate_t;
3183
3184     typedef __mbstate_t mbstate_t;
3185     extern int iswblank(wint_t);
3186     extern wint_t towlower(wint_t);
3187     extern wint_t towupper(wint_t);
3188     extern wctrans_t wctrans(const char *);
3189     extern int iswalnum(wint_t);
3190     extern int iswalpha(wint_t);
3191     extern int iswcntrl(wint_t);
3192     extern int iswctype(wint_t, wctype_t);
3193     extern int iswdigit(wint_t);
3194     extern int iswgraph(wint_t);
3195     extern int iswlower(wint_t);
3196     extern int iswprint(wint_t);
3197     extern int iswpunct(wint_t);
3198     extern int iswspace(wint_t);
3199     extern int iswupper(wint_t);
3200     extern int iswdxdigit(wint_t);
3201     extern wctype_t wctype(const char *);
3202     extern wint_t towctrans(wint_t, wctrans_t);

```

### 13.4.80 wordexp.h

```

3203
3204     enum {
3205         WRDE_DOOFFS, WRDE_APPEND, WRDE_NOCMD, WRDE_REUSE, WRDE_SHOWERR,
3206         WRDE_UNDEF, __WRDE_FLAGS
3207     };
3208
3209     typedef struct {
3210         int we_wordc;
3211         char **we_wordv;
3212         int we_offs;
3213     } wordexp_t;
3214
3215     enum {
3216         WRDE_NOSYS, WRDE_NOSPACE, WRDE_BADCHAR, WRDE_BADVAL, WRDE_CMDSUB,
3217         WRDE_SYNTAX
3218     };
3219     extern int wordexp(const char *, wordexp_t *, int);
3220     extern void wordfree(wordexp_t *);

```

## 13.5 Interface Definitions for libc

3221        The interfaces defined on the following pages are included in libc and are defined by  
 3222        this specification. Unless otherwise noted, these interfaces shall be included in the  
 3223        source standard.

3224        Other interfaces listed in Section 13.3 shall behave as described in the referenced  
 3225        base document.

### **\_IO\_feof**

#### **Name**

3226        `_IO_feof` — alias for `feof`

#### **Synopsis**

3227        `int _IO_feof(_IO_FILE * __fp);`

#### **Description**

3228        `_IO_feof()` tests the end-of-file indicator for the stream pointed to by `__fp`,  
 3229        returning a non-zero value if it is set.

3230        `_IO_feof()` is not in the source standard; it is only in the binary standard.

### **\_IO\_getc**

#### **Name**

3231        `_IO_getc` — alias for `getc`

#### **Synopsis**

3232        `int _IO_getc(_IO_FILE * __fp);`

#### **Description**

3233        `_IO_getc()` reads the next character from `__fp` and returns it as an unsigned char  
 3234        cast to an int, or EOF on end-of-file or error.

3235        `_IO_getc()` is not in the source standard; it is only in the binary standard.

### **\_IO\_putc**

#### **Name**

3236        `_IO_putc` — alias for `putc`

#### **Synopsis**

3237        `int _IO_putc(int __c, _IO_FILE * __fp);`

#### **Description**

3238        `_IO_putc()` writes the character `__c`, cast to an unsigned char, to `__fp`.

3239        `_IO_putc()` is not in the source standard; it is only in the binary standard.

## \_IO\_puts

### **Name**

3240    `_IO_puts` – alias for `puts`

### **Synopsis**

3241    `int _IO_puts(const char * __c);`

### **Description**

3242    `_IO_puts()` writes the string `__s` and a trailing newline to `stdout`.

3243    `_IO_puts()` is not in the source standard; it is only in the binary standard.

## \_assert\_fail

### **Name**

3244    `_assert_fail` – abort the program after false assertion

### **Synopsis**

3245    `void _assert_fail(const char * assertion, const char * file, unsigned int`  
 3246    `line, const char * function);`

### **Description**

3247    The `_assert_fail()` function is used to implement the `assert()` interface of ISO  
 3248    POSIX (2003). The `_assert_fail()` function shall print the given `file` filename,  
 3249    `line` line number, `function` function name and a message on the standard error  
 3250    stream in an unspecified format, and abort program execution via the `abort()`  
 3251    function. For example:

3252    a.c:10: foobar: Assertion a == b failed.

3253    If `function` is NULL, `_assert_fail()` shall omit information about the function.

3254    `assertion`, `file`, and `line` shall be non-NUL.

3255    The `_assert_fail()` function is not in the source standard; it is only in the binary  
 3256    standard. The `assert()` interface is not in the binary standard; it is only in the  
 3257    source standard. The `assert()` may be implemented as a macro.

**\_\_ctype\_b\_loc****Name**

3258       `__ctype_b_loc` — accessor function for `__ctype_b` array for ctype functions

**Synopsis**

3259       `#include <ctype.h>`  
 3260       `const unsigned short * * __ctype_b_loc (void);`

**Description**

3261       The `__ctype_b_loc()` function shall return a pointer into an array of characters in  
 3262       the current locale that contains characteristics for each character in the current  
 3263       character set. The array shall contain a total of 384 characters, and can be indexed  
 3264       with any signed or unsigned char (i.e. with an index value between -128 and 255). If  
 3265       the application is multithreaded, the array shall be local to the current thread.

3266       This interface is not in the source standard; it is only in the binary standard.

**Return Value**

3267       The `__ctype_b_loc()` function shall return a pointer to the array of characters to be  
 3268       used for the `ctype()` family of functions (see `<ctype.h>`).

**\_\_ctype\_get\_mb\_cur\_max****Name**

3269       `__ctype_get_mb_cur_max` — maximum length of a multibyte character in the  
 3270       current locale

**Synopsis**

3271       `size_t __ctype_get_mb_cur_max(void);`

**Description**

3272       `__ctype_get_mb_cur_max()` returns the maximum length of a multibyte character  
 3273       in the current locale.

3274       `__ctype_get_mb_cur_max()` is not in the source standard; it is only in the binary  
 3275       standard.

## **\_\_ctype\_tolower\_loc**

### **Name**

3276        \_\_ctype\_tolower\_loc — accessor function for \_\_ctype\_b\_tolower array for  
 3277        ctype tolower() function

### **Synopsis**

3278        #include <ctype.h>  
 3279        int32\_t \* \* \_\_ctype\_tolower\_loc(void);

### **Description**

3280        The \_\_ctype\_tolower\_loc() function shall return a pointer into an array of  
 3281        characters in the current locale that contains lower case equivalents for each  
 3282        character in the current character set. The array shall contain a total of 384 characters,  
 3283        and can be indexed with any signed or unsigned char (i.e. with an index value  
 3284        between -128 and 255). If the application is multithreaded, the array shall be local to  
 3285        the current thread.

3286        This interface is not in the source standard; it is only in the binary standard.

### **Return Value**

3287        The \_\_ctype\_tolower\_loc() function shall return a pointer to the array of  
 3288        characters to be used for the ctype() family of functions (see <ctype.h>).

## **\_\_ctype\_toupper\_loc**

### **Name**

3289        \_\_ctype\_toupper\_loc — accessor function for \_\_ctype\_b\_toupper() array for  
 3290        ctype toupper() function

### **Synopsis**

3291        #include <ctype.h>  
 3292        int32\_t \* \* \_\_ctype\_toupper\_loc(void);

### **Description**

3293        The \_\_ctype\_toupper\_loc() function shall return a pointer into an array of  
 3294        characters in the current locale that contains upper case equivalents for each  
 3295        character in the current character set. The array shall contain a total of 384 characters,  
 3296        and can be indexed with any signed or unsigned char (i.e. with an index value  
 3297        between -128 and 255). If the application is multithreaded, the array shall be local to  
 3298        the current thread.

3299        This interface is not in the source standard; it is only in the binary standard.

### **Return Value**

3300        The \_\_ctype\_toupper\_loc() function shall return a pointer to the array of  
 3301        characters to be used for the ctype() family of functions (see <ctype.h>).

## \_\_cxa\_atexit

### **Name**

3302       `__cxa_atexit` — register a function to be called by exit or when a shared library is  
 3303       unloaded

### **Synopsis**

3304       `int __cxa_atexit(void (*func) (void *), void * arg, void * dso_handle);`

### **Description**

3305       As described in the Itanium C++ ABI, `__cxa_atexit()` registers a destructor  
 3306       function to be called by `exit()` or when a shared library is unloaded. When a shared  
 3307       library is unloaded, any destructor function associated with that shared library,  
 3308       identified by `dso_handle`, shall be called with the single argument `arg`, and then  
 3309       that function shall be removed, or marked as complete, from the list of functions to  
 3310       run at `exit()`. On a call to `exit()`, any remaining functions registered shall be  
 3311       called with the single argument `arg`. Destructor functions shall always be called in  
 3312       the reverse order to their registration (i.e. the most recently registered function shall  
 3313       be called first),

3314       The `__cxa_atexit()` function is used to implement `atexit()`, as described in ISO  
 3315       POSIX (2003). Calling `atexit(func)` from the statically linked part of an application  
 3316       shall be equivalent to `__cxa_atexit(func, NULL, NULL)`.

3317       `__cxa_atexit()` is not in the source standard; it is only in the binary standard.

3318       **Note:** `atexit()` is not in the binary standard; it is only in the source standard.

## \_\_daylight

### **Name**

3319       `__daylight` — daylight savings time flag

### **Synopsis**

3320       `int __daylight;`

### **Description**

3321       The integer variable `__daylight` shall implement the daylight savings time flag  
 3322       `daylight` as specified in the ISO POSIX (2003) header file `<time.h>`.

3323       `__daylight` is not in the source standard; it is only in the binary standard. `daylight`  
 3324       is not in the binary standard; it is only in the source standard.

## \_\_environ

### **Name**

3325        \_\_environ – alias for environ - user environment

### **Synopsis**

3326        `extern char * * __environ;`

### **Description**

3327        \_\_environ is an alias for environ - user environment.

3328        \_\_environ has the same specification as environ.

3329        \_\_environ is not in the source standard; it is only in the binary standard.

## \_\_errno\_location

### **Name**

3330        \_\_errno\_location – address of errno variable

### **Synopsis**

3331        `int * __errno_location(void);`

### **Description**

3332        The \_\_errno\_location() function shall return the address of the errno variable for  
3333        the current thread.

3334        \_\_errno\_location() is not in the source standard; it is only in the binary standard.

## \_\_fpending

### **Name**

3335        \_\_fpending – returns in bytes the amount of output pending on a stream

### **Synopsis**

3336        `size_t __fpending(FILE * stream);`

### **Description**

3337        \_\_fpending() returns the amount of output in bytes pending on a stream.

3338        \_\_fpending() is not in the source standard; it is only in the binary standard.

**\_\_getpagesize****Name**

3339       `__getpagesize` — alias for `getpagesize` - get current page size

**Synopsis**

3340       `int __getpagesize(void);`

**Description**

3341       `__getpagesize()` is an alias for `getpagesize()` - get current page size.

3342       `__getpagesize()` has the same specification as `getpagesize()`.

3343       `__getpagesize()` is not in the source standard; it is only in the binary standard.

**\_\_getpgid****Name**

3344       `__getpgid` — get the process group id

**Synopsis**

3345       `pid_t __getpgid(pid_t pid);`

**Description**

3346       `__getpgid()` has the same specification as `getpgid()`.

3347       `__getpgid()` is not in the source standard; it is only in the binary standard.

**\_\_h\_errno\_location****Name**

3348       `__h_errno_location` — address of `h_errno` variable

**Synopsis**

3349       `int * __h_errno_location(void);`

**Description**

3350       `__h_errno_location()` returns the address of the `h_errno` variable, where `h_errno` is as specified in ISO POSIX (2003).

3352       `__h_errno_location()` is not in the source standard; it is only in the binary standard. Note that `h_errno` itself is only in the source standard; it is not in the binary standard.

## \_\_isinf

### **Name**

3355       `__isinf` — test for infinity

### **Synopsis**

3356       `int __isinf(double arg);`

### **Description**

3357       `__isinf()` has the same specification as `isinf()` in ISO POSIX (2003), except that  
3358       the argument type for `__isinf()` is known to be double.

3359       `__isinf()` is not in the source standard; it is only in the binary standard.

## \_\_isinff

### **Name**

3360       `__isinff` — test for infinity

### **Synopsis**

3361       `int __isinff(float arg);`

### **Description**

3362       `__isinff()` has the same specification as `isinf()` in ISO POSIX (2003) except that  
3363       the argument type for `__isinff()` is known to be float.

3364       `__isinff()` is not in the source standard; it is only in the binary standard.

## \_\_isinfl

### **Name**

3365       `__isinfl` — test for infinity

### **Synopsis**

3366       `int __isinfl(long double arg);`

### **Description**

3367       `__isinfl()` has the same specification as `isinf()` in the ISO POSIX (2003), except  
3368       that the argument type for `__isinfl()` is known to be long double.

3369       `__isinfl()` is not in the source standard; it is only in the binary standard.

**\_\_isnan****Name**

3370       `__isnan — test for infinity`

**Synopsis**

3371       `int __isnan(double arg);`

**Description**

3372       `__isnan()` has the same specification as `isnan()` in ISO POSIX (2003), except that  
3373       the argument type for `__isnan()` is known to be double.

3374       `__isnan()` is not in the source standard; it is only in the binary standard.

**\_\_isnanf****Name**

3375       `__isnanf — test for infinity`

**Synopsis**

3376       `int __isnanf(float arg);`

**Description**

3377       `__isnanf()` has the same specification as `isnan()` in ISO POSIX (2003), except that  
3378       the argument type for `__isnanf()` is known to be float.

3379       `__isnanf()` is not in the source standard; it is only in the binary standard.

**\_\_isnanl****Name**

3380       `__isnanl — test for infinity`

**Synopsis**

3381       `int __isnanl(long double arg);`

**Description**

3382       `__isnanl()` has the same specification as `isnan()` in ISO POSIX (2003), except that  
3383       the argument type for `__isnanl()` is known to be long double.

3384       `__isnanl()` is not in the source standard; it is only in the binary standard.

## \_\_libc\_current\_sigrtmax

### **Name**

3385        \_\_libc\_current\_sigrtmax — return number of available real-time signal with  
3386        lowest priority

### **Synopsis**

3387        int \_\_libc\_current\_sigrtmax(void);

### **Description**

3388        \_\_libc\_current\_sigrtmax() returns the number of an available real-time signal  
3389        with the lowest priority.

3390        \_\_libc\_current\_sigrtmax() is not in the source standard; it is only in the binary  
3391        standard.

## \_\_libc\_current\_sigrtmin

### **Name**

3392        \_\_libc\_current\_sigrtmin — return number of available real-time signal with  
3393        highest priority

### **Synopsis**

3394        int \_\_libc\_current\_sigrtmin(void);

### **Description**

3395        \_\_libc\_current\_sigrtmin() returns the number of an available real-time signal  
3396        with the highest priority.

3397        \_\_libc\_current\_sigrtmin() is not in the source standard; it is only in the binary  
3398        standard.

## \_\_libc\_start\_main

### Name

3399        \_\_libc\_start\_main — initialization routine

### Synopsis

3400        int \_\_libc\_start\_main(int \*(main) (int, char \*\*, char \*\*), int argc, char  
 3401                  \* \* *ubp\_av*, void (\*init) (void), void (\*fini) (void), void (\*rtld\_fini)  
 3402                  (void), void (\*stack\_end));

### Description

3403        The `__libc_start_main()` function shall perform any necessary initialization of the  
 3404                  execution environment, call the `main` function with appropriate arguments, and  
 3405                  handle the return from `main()`. If the `main()` function returns, the return value shall  
 3406                  be passed to the `exit()` function.

3407        **Note:** While this specification is intended to be implementation independent, process  
 3408                  and library initialization may include:

- 3409                  • performing any necessary security checks if the effective user ID is not the same as the  
   3410                      real user ID.
- 3411                  • initialize the threading subsystem.
- 3412                  • registering the `rtld_fini` to release resources when this dynamic shared object exits  
   3413                      (or is unloaded).
- 3414                  • registering the `fini` handler to run at program exit.
- 3415                  • calling the initializer function (`*init()`).
- 3416                  • calling `main()` with appropriate arguments.
- 3417                  • calling `exit()` with the return value from `main()`.

3418        This list is an example only.

3419        `__libc_start_main()` is not in the source standard; it is only in the binary  
 3420                  standard.

### See Also

3421        The section on Process Initialization in each of the architecture specific supplements.

## \_\_lxstat

### Name

3422        \_\_lxstat — inline wrapper around call to `lxstat`

### Synopsis

3423        #include <ctype.h>  
 3424        int \_\_lxstat(int *version*, char \* *path*, struct stat *statbuf*);

### Description

3425        `__lxstat()` is an inline wrapper around call to `lxstat()`.

3426        `__lxstat()` is not in the source standard; it is only in the binary standard.

## \_\_mempcpy

### **Name**

3427       `__mempcpy` — copy given number of bytes of source to destination

### **Synopsis**

3428       `#include <string.h>`  
 3429       `ptr_t __mempcpy(ptr_t restrict dest, const ptr_t restrict src, size_t n);`

### **Description**

3430       `__mempcpy()` copies *n* bytes of source to destination, returning pointer to bytes after  
 3431       the last written byte.  
 3432       `__mempcpy()` is not in the source standard; it is only in the binary standard.

## \_\_rawmemchr

### **Name**

3433       `__rawmemchr` — scan memory

### **Synopsis**

3434       `#include <string.h>`  
 3435       `ptr_t __rawmemchr(const ptr_t s, int c);`

### **Description**

3436       `__rawmemchr()` searches in *s* for *c*.  
 3437       `__rawmemchr()` is a weak alias to `rawmemchr()`. It is similar to `memchr()`, but it has  
 3438       no length limit.  
 3439       `__rawmemchr()` is not in the source standard; it is only in the binary standard.

## \_\_register\_atfork

### **Name**

3440       `__register_atfork` — alias for `register_atfork`

### **Synopsis**

3441       `int __register_atfork(void (*prepare) (void), void (*parent) (void), void`  
 3442       `(*child) (void), void * __dso_handle);`

### **Description**

3443       `__register_atfork()` implements `pthread_atfork()` as specified in ISO POSIX  
 3444       (2003). The additional parameter `__dso_handle` allows a shared object to pass in its  
 3445       handle so that functions registered by `__register_atfork()` can be unregistered by  
 3446       the runtime when the shared object is unloaded.

**\_\_sigsetjmp****Name**

3447        \_\_sigsetjmp — save stack context for non-local goto

**Synopsis**

3448        int \_\_sigsetjmp( jmp\_buf env, int savemask);

**Description**

3449        \_\_sigsetjmp( ) has the same behavior as `sigsetjmp( )` as specified by ISO POSIX  
3450        (2003).

3451        \_\_sigsetjmp( ) is not in the source standard; it is only in the binary standard.

**\_\_stpcpy****Name**

3452        \_\_stpcpy — alias for `stpcpy`

**Synopsis**

3453        #include <string.h>  
3454        char \* \_\_stpcpy(char \* dest, const char \* src);

**Description**

3455        The `__stpcpy()` function has the same specification as the `stpcpy()`.

3456        `__stpcpy()` is not in the source standard; it is only in the binary standard.

**\_\_strdup****Name**

3457        \_\_strdup — alias for `strdup`

**Synopsis**

3458        char \* \_\_strdup(const char string);

**Description**

3459        `__strdup()` has the same specification as `strdup()`.

3460        `__strdup()` is not in the source standard; it is only in the binary standard.

## \_\_strtod\_internal

### **Name**

3461        \_\_strtod\_internal — underlying function for strtod

### **Synopsis**

3462        double \_\_strtod\_internal(const char \* \_\_nptr, char \* \* \_\_endptr, int \_\_group);

### **Description**

3463        \_\_group shall be 0 or the behavior of \_\_strtod\_internal() is undefined.

3464        \_\_strtod\_internal(\_\_nptr, \_\_endptr, 0)() has the same specification as  
3465        strtod(\_\_nptr, \_\_endptr)().

3466        \_\_strtod\_internal() is not in the source standard; it is only in the binary  
3467        standard.

## \_\_strtof\_internal

### **Name**

3468        \_\_strtof\_internal — underlying function for strtof

### **Synopsis**

3469        float \_\_strtof\_internal(const char \* \_\_nptr, char \* \* \_\_endptr, int \_\_group);

### **Description**

3470        \_\_group shall be 0 or the behavior of \_\_strtof\_internal() is undefined.

3471        \_\_strtof\_internal(\_\_nptr, \_\_endptr, 0)() has the same specification as  
3472        strtof(\_\_nptr, \_\_endptr)().

3473        \_\_strtof\_internal() is not in the source standard; it is only in the binary  
3474        standard.

## \_\_strtok\_r

### **Name**

3475        \_\_strtok\_r — alias for strtok\_r

### **Synopsis**

3476        char \* \_\_strtok\_r(char \* restrict s, const char \* restrict delim, char \*  
3477        \* restrict save\_ptr);

### **Description**

3478        \_\_strtok\_r() has the same specification as strtok\_r().

3479        \_\_strtok\_r() is not in the source standard; it is only in the binary standard.

**\_\_strtol\_internal****Name**

3480       *\_\_strtol\_internal* — alias for *strtol*

**Synopsis**

3481       long int *\_\_strtol\_internal*(const char \* *\_nptr*, char \* \* *\_endptr*, int *\_base*,  
3482       int *\_group*);

**Description**

3483       *\_group* shall be 0 or the behavior of *\_\_strtol\_internal()* is undefined.

3484       *\_\_strtol\_internal*(*\_nptr*, *\_endptr*, *\_base*, 0) has the same specification as  
3485       *strtol*(*\_nptr*, *\_endptr*, *\_base*).

3486       *\_\_strtol\_internal()* is not in the source standard; it is only in the binary  
3487       standard.

**\_\_strtold\_internal****Name**

3488       *\_\_strtold\_internal* — underlying function for *strtold*

**Synopsis**

3489       long double *\_\_strtold\_internal*(const char \* *\_nptr*, char \* \* *\_endptr*, int  
3490       *\_group*);

**Description**

3491       *\_group* shall be 0 or the behavior of *\_\_strtold\_internal()* is undefined.

3492       *\_\_strtold\_internal*(*\_nptr*, *\_endptr*, 0) has the same specification as  
3493       *strtold*(*\_nptr*, *\_endptr*).

3494       *\_\_strtold\_internal()* is not in the source standard; it is only in the binary  
3495       standard.

## \_\_strtoll\_internal

### **Name**

3496        \_\_strtoll\_internal — underlying function for strtoll

### **Synopsis**

3497        long long \_\_strtoll\_internal(const char \* \_\_nptr, char \* \* \_\_endptr, int \_\_base,  
3498            int \_\_group);

### **Description**

3499        \_\_group shall be 0 or the behavior of \_\_strtoll\_internal() is undefined.

3500        \_\_strtoll\_internal(\_\_nptr, \_\_endptr, \_\_base, 0) has the same specification as  
3501            strtoll(\_\_nptr, \_\_endptr, \_\_base).

3502        \_\_strtoll\_internal() is not in the source standard; it is only in the binary  
3503            standard.

## \_\_strtoul\_internal

### **Name**

3504        \_\_strtoul\_internal — underlying function for strtoul

### **Synopsis**

3505        unsigned long int \_\_strtoul\_internal(const char \* \_\_nptr, char \* \* \_\_endptr,  
3506            int \_\_base, int \_\_group);

### **Description**

3507        \_\_group shall be 0 or the behavior of \_\_strtoul\_internal() is undefined.

3508        \_\_strtoul\_internal(\_\_nptr, \_\_endptr, \_\_base, 0) has the same specification as  
3509            strtoul(\_\_nptr, \_\_endptr, \_\_base).

3510        \_\_strtoul\_internal() is not in the source standard; it is only in the binary  
3511            standard.

## \_\_strtoull\_internal

### **Name**

3512        \_\_strtoull\_internal — underlying function for strtoull

### **Synopsis**

3513        unsigned long long \_\_strtoull\_internal(const char \* \_\_nptr, char \* \* \_\_endptr,  
3514                    int \_\_base, int \_\_group);

### **Description**

3515        \_\_group shall be 0 or the behavior of \_\_strtoull\_internal() is undefined.

3516        \_\_strtoull\_internal(\_\_nptr, \_\_endptr, \_\_base, 0) has the same specification as  
3517                    strtoull(\_\_nptr, \_\_endptr, \_\_base).

3518        \_\_strtoull\_internal() is not in the source standard; it is only in the binary  
3519                    standard.

## \_\_sysconf

### **Name**

3520        \_\_sysconf — get configuration information at runtime

### **Synopsis**

3521        #include <unistd.h>  
3522        long \_\_sysconf(int name);

### **Description**

3523        \_\_sysconf() gets configuration information at runtime.

3524        \_\_sysconf() is weak alias to sysconf().

3525        \_\_sysconf() has the same specification as sysconf().

3526        \_\_sysconf() is not in the source standard; it is only in the binary standard.

## \_\_sysv\_signal

### **Name**

3527        \_\_sysv\_signal — signal handling

### **Synopsis**

3528        \_\_sighandler\_t \_\_sysv\_signal(int sig, \_\_sighandler\_t handler);

### **Description**

3529        \_\_sysv\_signal() has the same behavior as signal() as specified by ISO POSIX  
3530                    (2003).

3531        \_\_sysv\_signal() is not in the source standard; it is only in the binary standard.

## \_\_timezone

### **Name**

3532 – global variable containing timezone

### **Synopsis**

3533 `long int __timezone;`

### **Description**

3534 `__timezone( )` has the same specification as `timezone( )` in the *ISO POSIX (2003)*

## \_\_tzname

### **Name**

3535 – global variable containing the timezone

### **Synopsis**

3536 `char * __tzname[ 2 ];`

### **Description**

3537 `__tzname` has the same specification as `tzname` in the *ISO POSIX (2003)*.

3538 Note that the array size of 2 is explicit in the *ISO POSIX (2003)*, but not in the *SUSv2*.

## \_\_wcstod\_internal

### **Name**

3539 `__wcstod_internal` – underlying function for `wcstod`

### **Synopsis**

3540 `double __wcstod_internal(const wchar_t * nptr, wchar_t * * endptr, int group);`

### **Description**

3541 `group` shall be 0 or the behavior of `__wcstod_internal()` is undefined.

3542 `__wcstod_internal(nptr, endptr, 0)` shall behave as `wcstod(nptr, endptr)` as  
3543 specified by ISO POSIX (2003).

3544 `__wcstod_internal()` is not in the source standard; it is only in the binary  
3545 standard.

**\_\_wcstof\_internal****Name**

3546       `__wcstof_internal` — underlying function for `wcstof`

**Synopsis**

3547       `float __wcstof_internal(const wchar_t * nptr, wchar_t * * endptr, int group);`

**Description**

3548       `group` shall be 0 or the behavior of `__wcstof_internal()` is undefined.

3549       `__wcstof_internal(nptr, endptr, 0)` shall behave as `wcstof(nptr, endptr)` as  
3550       specified in ISO POSIX (2003).

3551       `__wcstof_internal()` is not in the source standard; it is only in the binary  
3552       standard.

**\_\_wcstol\_internal****Name**

3553       `__wcstol_internal` — underlying function for `wcstol`

**Synopsis**

3554       `long __wcstol_internal(const wchar_t * nptr, wchar_t * * endptr, int base,`  
3555       `int group);`

**Description**

3556       `group` shall be 0 or the behavior of `__wcstol_internal()` is undefined.

3557       `__wcstol_internal(nptr, endptr, base, 0)` shall behave as `wcstol(nptr, endptr,`  
3558       `base)` as specified by ISO POSIX (2003).

3559       `__wcstol_internal()` is not in the source standard; it is only in the binary  
3560       standard.

**\_\_wcstold\_internal****Name**

3561       `__wcstold_internal` — underlying function for `wcstold`

**Synopsis**

3562       `long double __wcstold_internal(const wchar_t * nptr, wchar_t * * endptr, int`  
3563       `group);`

**Description**

3564       `group` shall be 0 or the behavior of `__wcstold_internal()` is undefined.

3565       `__wcstold_internal(nptr, endptr, 0)` shall behave as `wcstold(nptr, endptr)` as  
3566       specified by ISO POSIX (2003).

3567       `__wcstold_internal()` is not in the source standard; it is only in the binary  
3568       standard.

## wcstoul\_internal

### Name

3569       `__wcstoul_internal` – underlying function for `wcstoul`

### Synopsis

3570       `unsigned long __wcstoul_internal(const wchar_t * restrict nptr, wchar_t *`  
 3571       `* restrict endptr, int base, int group);`

### Description

3572       `group` shall be 0 or the behavior of `__wcstoul_internal()` is undefined.

3573       `__wcstoul_internal(nptr, endptr, base, 0)()` shall behave as `wcstoul(nptr,`  
 3574       `endptr, base)()` as specified by ISO POSIX (2003).

3575       `__wcstoul_internal()` is not in the source standard; it is only in the binary  
 3576       standard.

## xmknod

### Name

3577       `__xmknod` – make block or character special file

### Synopsis

3578       `int __xmknod(int ver, const char * path, mode_t mode, dev_t * dev);`

### Description

3579       The `__xmknod()` function shall implement the `mknod()` interface from ISO POSIX  
 3580       (2003).

3581       The value of `ver` shall be 1 or the behavior of `__xmknod()` is undefined.

3582       `__xmknod(1, path, mode, dev)` shall behave as `mknod(path, mode, dev)` as specified  
 3583       by ISO POSIX (2003).

3584       The `__xmknod()` function is not in the source standard; it is only in the binary  
 3585       standard.

3586       **Note:** The `mknod()` function is not in the binary standard; it is only in the source  
 3587       standard.

## xstat

### Name

3588       `__xstat` – get File Status

### Synopsis

3589       `#include <sys/stat.h>`

```
3590     #include <unistd.h>
3591     int __xstat(int ver, const char * path, struct stat * stat_buf);
3592     int __lxstat(int ver, const char * path, struct stat * stat_buf);
3593     int __fxstat(int ver, int fildes, struct stat * stat_buf);
```

## Description

3594 The functions `__xstat()`, `__lxstat()`, and `__fxstat()` shall implement the ISO  
 3595 POSIX (2003) functions `stat()`, `lstat()`, and `fstat()` respectively.  
 3596 `ver` shall be 3 or the behavior of these functions is undefined.  
 3597 `__xstat(3, path, stat_buf)` shall implement `stat(path, stat_buf)` as specified by  
 3598 ISO POSIX (2003).  
 3599 `__lxstat(3, path, stat_buf)` shall implement `lstat(path, stat_buf)` as specified  
 3600 by ISO POSIX (2003).  
 3601 `__fxstat(3, fildes, stat_buf)` shall implement `fstat(fildes, stat_buf)` as  
 3602 specified by ISO POSIX (2003).  
 3603 `__xstat()`, `__lxstat()`, and `__fxstat()` are not in the source standard; they are  
 3604 only in the binary standard.  
 3605 `stat()`, `lstat()`, and `fstat()` are not in the binary standard; they are only in the  
 3606 source standard.

## **\_\_xstat64**

### Name

3607 `__xstat64` — get File Status

### Synopsis

```
3608 #define _LARGEFILE_SOURCE 1
3609 #include <sys/stat.h>
3610 #include <unistd.h>
3611 int __xstat64(int ver, const char * path, struct stat64 * stat_buf);
3612 int __lxstat64(int ver, const char * path, struct stat64 * stat_buf);
3613 int __fxstat64(int ver, int fildes, struct stat64 * stat_buf);
```

## Description

3614 The functions `__xstat64()`, `__lxstat64()`, and `__fxstat64()` shall implement the  
 3615 Large File Support functions `stat64()`, `lstat64()`, and `fstat64()` respectively.  
 3616 `ver` shall be 3 or the behavior of these functions is undefined.  
 3617 `__xstat64(3, path, stat_buf)` shall behave as `stat(path, stat_buf)` as specified  
 3618 by Large File Support.  
 3619 `__lxstat64(3, path, stat_buf)` shall behave as `lstat(path, stat_buf)` as specified  
 3620 by Large File Support.  
 3621 `__fxstat64(3, fildes, stat_buf)` shall behave as `fstat(fildes, stat_buf)` as  
 3622 specified by Large File Support.  
 3623 `__xstat64()`, `__lxstat64()`, and `__fxstat64()` are not in the source standard;  
 3624 they are only in the binary standard.  
 3625 `stat64()`, `lstat64()`, and `fstat64()` are not in the binary standard; they are only  
 3626 in the source standard.

**\_environ****Name**

3627        \_environ — alias for environ - user environment

**Synopsis**

3628        extern char \* \*\_environ;

**Description**

3629        \_environ is an alias for environ - user environment.

**\_nl\_msg\_cat\_cntr****Name**

3630        \_nl\_msg\_cat\_cntr — new catalog load counter

**Synopsis**

3631        #include <libintl.h>  
 3632  
 3633        extern int \_nl\_msg\_cat\_cntr;

**Description**

3634        The global variable \_nl\_msg\_cat\_cntr is incremented each time a new catalog is  
 3635        loaded. This variable is only in the binary standard; it is not in the source standard.

**\_sys\_errlist****Name**

3636        \_sys\_errlist — array containing the "C" locale strings used by strerror()

**Synopsis**

3637        #include <stdio.h>  
 3638  
 3639        extern const char \*const \_sys\_errlist[ ];

**Description**

3640        \_sys\_errlist is an array containing the "C" locale strings used by strerror(). This  
 3641        normally should not be used directly. strerror() provides all of the needed  
 3642        functionality.

**\_sys\_siglist****Name**

3643        \_sys\_siglist — array containing the names of the signal names

**Synopsis**

3644        #include <signal.h>  
 3645

```
3646     extern const char *const _sys_siglist[NSIG];
```

## Description

3647     `_sys_siglist` is an array containing the names of the signal names.

3648     The `_sys_siglist` array is only in the binary standard; it is not in the source  
3649     standard. Applications wishing to access the names of signals should use the  
3650     `strsignal()` function.

## acct

### Name

3651     `acct` — switch process accounting on or off

### Synopsis

```
3652 #include <dirent.h>
3653 int acct(const char * filename);
```

### Description

3654     When `filename` is the name of an existing file, `acct()` turns accounting on and  
3655     appends a record to `filename` for each terminating process. When `filename` is `NULL`,  
3656     `acct()` turns accounting off.

### Return Value

3657     On success, 0 is returned. On error, -1 is returned and the global variable `errno` is set  
3658     appropriately.

### Errors

3659     ENOSYS

3660         BSD process accounting has not been enabled when the operating system kernel  
3661         was compiled. The kernel configuration parameter controlling this feature is  
3662         `CONFIG_BSD_PROCESS_ACCT`.

3663     ENOMEM

3664         Out of memory.

3665     EPERM

3666         The calling process has no permission to enable process accounting.

3667     EACCES

3668         `filename` is not a regular file.

3669     EIO

3670         Error writing to the `filename`.

3671     EUSERS

3672         There are no more free file structures or we run out of memory.

## adjtime

### Name

3673 adjtime — correct the time to allow synchronization of the system clock

### Synopsis

```
3674 #include <time.h>
3675 int adjtime(const struct timeval * delta, struct timeval * olddelta);
```

### Description

3676 adjtime() makes small adjustments to the system time as returned by  
 3677 gettimeofday()(2), advancing or retarding it by the time specified by the timeval  
 3678 *delta*. If *delta* is negative, the clock is slowed down by incrementing it more slowly  
 3679 than normal until the correction is complete. If *delta* is positive, a larger increment  
 3680 than normal is used. The skew used to perform the correction is generally a fraction  
 3681 of one percent. Thus, the time is always a monotonically increasing function. A time  
 3682 correction from an earlier call to adjtime() may not be finished when adjtime() is  
 3683 called again. If *olddelta* is non-NULL, the structure pointed to will contain, upon  
 3684 return, the number of microseconds still to be corrected from the earlier call.

3685 adjtime() may be used by time servers that synchronize the clocks of computers in  
 3686 a local area network. Such time servers would slow down the clocks of some  
 3687 machines and speed up the clocks of others to bring them to the average network  
 3688 time.

3689 Appropriate privilege is required to adjust the system time.

### Return Value

3690 On success, 0 is returned. On error, -1 is returned and the global variable errno is set  
 3691 appropriately.

### Errors

3692 EFAULT

3693 An argument points outside the process's allocated address space.

3694 EPERM

3695 The process does not have appropriate privilege.

## asprintf

### Name

3696        `asprintf` — write formatted output to a dynamically allocated string

### Synopsis

3697        `#include <stdio.h>`  
3698        `int asprintf(char ** restrict ptr, const char * restrict format, ...);`

### Description

3699        The `asprintf()` function shall behave as `sprintf()`, except that the output string  
3700        shall be dynamically allocated space of sufficient length to hold the resulting string.  
3701        The address of this dynamically allocated string shall be stored in the location  
3702        referenced by `ptr`.

### Return Value

3703        Refer to `fprintf()`.

### Errors

3704        Refer to `fprintf()`.

## bind\_textdomain\_codeset

### Name

3705 bind\_textdomain\_codeset — specify encoding for message retrieval

### Synopsis

```
3706 #include <libintl.h>
3707 char * bind_textdomain_codeset (const char * domainname , const char *
3708 codeset );
```

### Description

3709 The `bind_textdomain_codeset()` function can be used to specify the output  
 3710 codeset for message catalogs for domain `domainname`. The `codeset` argument shall  
 3711 be a valid codeset name which can be used for the `iconv_open` function, or a null  
 3712 pointer. If the `codeset` argument is the null pointer, then function returns the  
 3713 currently selected codeset for the domain with the name `domainname`. It shall return  
 3714 a null pointer if no codeset has yet been selected.

3715 Each successive call to `bind_textdomain_codeset()` function overrides the  
 3716 settings made by the preceding call with the same `domainname`.

3717 The `bind_textdomain_codeset()` function shall return a pointer to a string  
 3718 containing the name of the selected codeset. The string shall be allocated internally  
 3719 in the function and shall not be changed or freed by the user.

3720 The `bind_textdomain_codeset()` function returns a pointer to a string containing  
 3721 the name of the selected codeset. The string is allocated internally in the function  
 3722 and shall not be changed by the user.

### Parameters

3723 `domainname`

3724 The `domainname` argument is applied to the currently active LC\_MESSAGE  
 3725 locale. It is equivalent in syntax and meaning to the `domainname` argument to  
 3726 `textdomain`, except that the selection of the domain is valid only for the  
 3727 duration of the call.

3728 `codeset`

3729 The name of the output codeset for the selected domain, or NULL to select the  
 3730 current codeset.

3731 If `domainname` is the null pointer, or is an empty string,  
 3732 `bind_textdomain_codeset()` shall fail, but need not set `errno`.

### Return Value

3733 Returns the currently selected codeset name. It returns a null pointer if no codeset  
 3734 has yet been selected.

### Errors

3735 ENOMEM

3736 Insufficient memory available to allocate return value.

## See Also

3737      gettext, dgettext, ngettext, dngettext, dcgettext, dcngettext, textdomain,  
3738      bindtextdomain

# bindresvport

## Name

3739      bindresvport — bind socket to privileged IP port

## Synopsis

```
3740 #include <sys/types.h>
3741 #include <rpc/rpc.h>
3742 int bindresvport(int sd, struct sockaddr_in * sin);
```

## Description

3743      If the process has appropriate privilege, the `bindresvport()` function shall bind a  
3744      socket to a privileged IP port.

## Return Value

3745      On success, 0 is returned. On error, -1 is returned and the global variable `errno` is set  
3746      appropriately.

## Errors

3747      EPERM  
3748      The process did not have appropriate privilege.  
3749      EPFNOSUPPORT  
3750      Address of *sin* did not match address family of *sd*.

## bindtextdomain

### Name

3751 bindtextdomain — specify the location of a message catalog

### Synopsis

```
3752 #include <libintl.h>
3753 char * bindtextdomain(const char * domainname, const char * dirname);
```

### Description

3754 The `bindtextdomain()` shall set the the base directory of the hierarchy containing  
 3755 message catalogs for a given message domain.

3756 The `bindtextdomain()` function specifies that the `domainname` message catalog can  
 3757 be found in the `dirname` directory hierarchy, rather than in the system default locale  
 3758 data base.

3759 If `dirname` is not `NULL`, the base directory for message catalogs belonging to domain  
 3760 `domainname` shall be set to `dirname`. If `dirname` is `NULL`, the base directory for  
 3761 message catalogs shall not be altered.

3762 The function shall make copies of the argument strings as needed.

3763 `dirname` can be an absolute or relative pathname.

3764 **Note:** Applications that wish to use `chdir()` should always use absolute pathnames to  
 3765 avoid inadvertently selecting the wrong or non-existent directory.

3766 If `domainname` is the null pointer, or is an empty string, `bindtextdomain()` shall fail,  
 3767 but need not set `errno`.

3768 The `bindtextdomain()` function shall return a pointer to a string containing the  
 3769 name of the selected directory. The string shall be allocated internally in the function  
 3770 and shall not be changed or freed by the user.

### Return Value

3771 On success, `bindtextdomain()` shall return a pointer to a string containing the  
 3772 directory pathname currently bound to the domain. On failure, a `NULL` pointer is  
 3773 returned, and the global variable `errno` may be set to indicate the error.

### Errors

3774 ENOMEM

3775 Insufficient memory was available.

### See Also

3776 `gettext`, `dgettext`, `ngettext`, `dngettext`, `dcgettext`, `dcngettext`, `textdomain`,  
 3777 `bind_textdomain_codeset`

## cfmakeraw

### Name

3778 `cfmakeraw` — get and set terminal attributes

### Synopsis

```
3779 #include <termios.h>
3780 void cfmakeraw(struct termios * termios_p);
```

### Description

3781 The `cfmakeraw()` function shall set the attributes of the `termios` structure  
 3782 referenced by `termios_p` as follows:

```
3783     termios_p->c_iflag &= ~(IGNBRK|BRKINT|PARMRK|ISTRIP
3784                               |INLCR|IGNCR|ICRNL|IXON);
3785
3786     termios_p->c_oflag &= ~OPOST;
3787
3788     termios_p->c_lflag &= ~(ECHO|ECHONL|ICANON|ISIG|IEXTEN);
3789
3790     termios_p->c_cflag &= ~(CSIZE|PARENNB);
3791
3792     termios_p->c_cflag |= CS8;
```

3793 `termios_p` shall point to a `termios` structure that contains the following members:

```
3794     tcflag_t c_iflag;          /* input modes */
3795     tcflag_t c_oflag;          /* output modes */
3796     tcflag_t c_cflag;          /* control modes */
3797     tcflag_t c_lflag;          /* local modes */
3798     cc_t c_cc[NCCS];          /* control chars */
```

## cfsetspeed

### Name

3799           `cfsetspeed — set terminal input and output data rate`

### Synopsis

3800           `#include <termios.h>`  
 3801           `int cfsetspeed(struct termios *t, speed_t speed);`

### Description

3802           `cfsetspeed()` sets the baud rate values in the `termios` structure. The effects of the  
 3803           function on the terminal as described below do not become effective, nor are all  
 3804           errors detected, until the `tcsetattr()` function is called. Certain values for baud  
 3805           rates set in `termios` and passed to `tcsetattr()` have special meanings.

### Getting and Setting the Baud Rate

3807           Input and output baud rates are found in the `termios` structure. The unsigned  
 3808           integer `speed_t` is typdef'd in the include file `termios.h`. The value of the integer  
 3809           corresponds directly to the baud rate being represented; however, the following  
 3810           symbolic values are defined.

```
3811           #define B0      0
3812           #define B50     50
3813           #define B75     75
3814           #define B110    110
3815           #define B134    134
3816           #define B150    150
3817           #define B200    200
3818           #define B300    300
3819           #define B600    600
3820           #define B1200   1200
3821           #define B1800   1800
3822           #define B2400   2400
3823           #define B4800   4800
3824           #define B9600   9600
3825           #define B19200  19200
3826           #define B38400  38400
3827           #ifndef __POSIX_SOURCE
3828           #define EXTA    19200
3829           #define EXTB    38400
3830           #endif /* __POSIX_SOURCE */
```

3831           `cfsetspeed()` sets both the input and output baud rates in the `termios` structure  
 3832           referenced by `t` to `speed`.

### Return Value

3833           On success, 0 is returned. On error, -1 is returned and the global variable `errno` is set  
 3834           appropriately.

### Errors

3835           EINVAL  
 3836           Invalid `speed` argument

## daemon

### Name

3837           daemon — run in the background

### Synopsis

3838           #include <unistd.h>  
3839           int daemon(int nochdir, int noclose);

### Description

3840           The `daemon()` function shall create a new process, detached from the controlling  
3841           terminal. If successful, the calling process shall exit and the new process shall  
3842           continue to execute the application in the background. If `nochdir` evaluates to true,  
3843           the current directory shall not be changed. Otherwise, `daemon()` shall change the  
3844           current working directory to the root ('/'). If `noclose` evaluates to true the standard  
3845           input, standard output, and standard error file descriptors shall not be altered.  
3846           Otherwise, `daemon()` shall close the standard input, standard output and standard  
3847           error file descriptors and reopen them attached to `/dev/null`.

### Return Value

3848           On error, -1 is returned, and the global variable `errno` is set to any of the errors  
3849           specified for the library functions `fork()` and `setsid()`.

## dcgettext

### Name

3850           dcgettext — perform domain and category specific lookup in message catalog

### Synopsis

3851           #include <libintl.h>

```
3852 #include <locale.h>
3853 char * dcgettext(const char * domainname, const char * msgid, int category);
```

## Description

3854 The `dcgettext()` function is a domain specified version of `gettext()`.

3855  
 3856 The `dcgettext()` function shall lookup the translation in the current locale of the  
 3857 message identified by `msgid` in the domain specified by `domainname` and in the  
 3858 locale category specified by `category`. If `domainname` is NULL, the current default  
 3859 domain shall be used. The `msgid` argument shall be a NULL-terminated string to be  
 3860 matched in the catalogue. `category` shall specify the locale category to be used for  
 3861 retrieving message strings. The `category` parameter shall be one of `LC_CTYPE`,  
 3862 `LC_COLLATE`, `LC_MESSAGES`, `LC_MONETARY`, `LC_NUMERIC`, or `LC_TIME`. The default  
 domain shall not be changed by a call to `dcgettext()`.

## Return Value

3863 If a translation was found in one of the specified catalogs, it shall be converted to the  
 3864 current locale's codeset and returned. The resulting NULL-terminated string shall be  
 3865 allocated by the `dcgettext` function, and must not be modified or freed. If no  
 3866 translation was found, or category was invalid, `msgid` shall be returned.

## Errors

3867 `dcgettext()` shall not modify the `errno` global variable.

## See Also

3868 `gettext`, `dgettext`, `ngettext`, `dngettext`, `dcngettext`, `textdomain`, `bindtextdomain`,  
 3869 `bind_textdomain_codeset`

## **dcngettext**

### Name

3870 `dcngettext` – perform domain and category specific lookup in message catalog  
 3871 with plural

### Synopsis

3872 `#include <libintl.h>`

```
3873     #include <locale.h>
3874     char * dcngettext(const char * domainname, const char * msgid1, const char
3875     * msgid2, unsigned long int n, int category);
```

## Description

3876 The `dcngettext()` function is a domain specific version of `gettext`, capable of  
 3877 returning either a singular or plural form of the message. The `dcngettext()`  
 3878 function shall lookup the translation in the current locale of the message identified  
 3879 by `msgid1` in the domain specified by `domainname` and in the locale category  
 3880 specified by `category`. If `domainname` is NULL, the current default domain shall be  
 3881 used. The `msgid1` argument shall be a NULL-terminated string to be matched in the  
 3882 catalogue. `category` shall specify the locale category to be used for retrieving  
 3883 message strings. The `category` parameter shall be one of `LC_CTYPE`, `LC_COLLATE`,  
 3884 `LC_MESSAGES`, `LC_MONETARY`, `LC_NUMERIC`, or `LC_TIME`. The default domain shall not  
 3885 be changed by a call to `dcngettext()`. If `n` is 1 then the singular version of the  
 3886 message is returned, otherwise one of the plural forms is returned, depending on the  
 3887 value of `n` and the current locale settings.

## Return Value

3888 If a translation corresponding to the value of `n` was found in one of the specified  
 3889 catalogs for `msgid1`, it shall be converted to the current locale's codeset and returned.  
 3890 The resulting NULL-terminated string shall be allocated by the `dcngettext()`  
 3891 function, and must not be modified or freed. If no translation was found, or  
 3892 `category` was invalid, `msgid1` shall be returned if `n` has the value 1, otherwise  
 3893 `msgid2` shall be returned.

## Errors

3894 `dcngettext()` shall not modify the `errno` global variable.

## See Also

3895 `gettext`, `dgettext`, `ngettext`, `dngettext`, `dcgettext`, `textdomain`, `bindtextdomain`,  
 3896 `bind_textdomain_codeset`

## dgettext

### Name

3897 dgettext — perform lookup in message catalog for the current LC\_MESSAGES  
 3898 locale

### Synopsis

```
3899 #include <libintl.h>
3900 char * dgettext(const char * domainname, const char * msgid);
```

### Description

3901 `dgettext()` is a domain specified version of `gettext()`.

3902 The `dgettext()` function shall search the currently selected message catalogs in the  
 3903 domain `domainname` for a string identified by the string `msgid`. If a string is located,  
 3904 that string shall be returned. The domain specified by `domainname` applies to the  
 3905 currently active LC\_MESSAGE locale. The default domain shall not be changed by a  
 3906 call to `dgettext()`.

3907 **Note:** The usage of `domainname` is equivalent in syntax and meaning to the  
 3908 `textdomain()` function's application of `domainname`, except that the selection of the  
 3909 domain in `dgettext()` is valid only for the duration of the call.

3910 The `dgettext()` function is equivalent to `dcgettext(domainname, msgid,`  
 3911 `LC_MESSAGES)`.

### Return Value

3912 On success of a `msgid` query, the translated NULL-terminated string is returned. On  
 3913 error, the original `msgid` is returned. The length of the string returned is  
 3914 undetermined until `dgettext()` is called.

### Errors

3915 `dgettext()` shall not modify the `errno` global variable.

### See Also

3916 `gettext`, `dgettext`, `ngettext`, `dngettext`, `dcgettext`, `dcngettext`, `textdomain`,  
 3917 `bindtextdomain`, `bind_textdomain_codeset`

**dnggettext****Name**

3918        `dnggettext` — perform lookup in message catalog for the current locale

**Synopsis**

```
3919 #include <libintl.h>
3920 char * dnggettext(const char * domainname, const char * msgid1, const char *
3921 msgid2, unsigned long int n);
```

**Description**

3922        `dnggettext()` shall be equivalent to a call to

3923        `dcngettext(domainname, msgid1, msgid2, n, LC_MESSAGES)`

3924        See `dcngettext()` for more information.

**See Also**

3925        `gettext`, `dgettext`, `ngettext`, `dcgettext`, `dcngettext`, `textdomain`, `bindtextdomain`,
 3926        `bind_textdomain_codeset`

**duplocale****Name**

3927        `duplocale` — provide new handle for selection of locale

**Synopsis**

```
3928 #include <locale.h>
3929 locale_t duplocale(locale_t locale);
```

**Description**

3930        The `duplocale()` function shall provide a new locale object based on the locale
 3931        object provided in `locale`, suitable for use in the `newlocale()` or `uselocale()`
 3932        functions. The new object may be released by calling `freelocale()`.

**Return Value**

3933        On success, the `duplocale()` function shall return a locale object. Otherwise, it shall
 3934        return `NULL`, and set `errno` to indicate the error.

**Errors**

3935        The `duplocale()` function shall fail if:

3936        ENOMEM

3937        Insufficient memory.

**See Also**

3938        `setlocale()`, `freelocale()`, `newlocale()`, `uselocale()`

## err

### Name

3939    `err` – display formatted error messages

### Synopsis

3940    `#include <err.h>`  
3941    `void err(int eval, const char * fmt, ...);`

### Description

3942    The `err()` function shall display a formatted error message on the standard error  
3943    stream. First, `err()` shall write the last component of the program name, a colon  
3944    character, and a space character. If `fmt` is non-NUL, it shall be used as a format  
3945    string for the `printf()` family of functions, and `err()` shall write the formatted  
3946    message, a colon character, and a space. Finally, the error message string affiliated  
3947    with the current value of the global variable `errno` shall be written, followed by a  
3948    newline character.

3949    The `err()` function shall not return, the program shall terminate with the exit value  
3950    of `eval`.

### See Also

3951    `error()`, `errx()`

### Return Value

3952    None.

### Errors

3953    None.

**error****Name**

3954       **error** — print error message

**Synopsis**

3955       **#include <err.h>**  
 3956       **void error(int exitstatus, int errnum, const char \* format, ...);**

**Description**

3957       **error()** shall print a message to standard error.

3958       **error()** shall build the message from the following elements in their specified  
 3959       order:

- 3960       1. the program name. If the application has provided a function named  
           3961       **error\_print\_progname()**, **error()** shall call this to supply the program  
           3962       name; otherwise, **error()** uses the content of the global variable  
           3963       **program\_name**.
- 3964       2. the colon and space characters, then the result of using the printf-style *format*  
           3965       and the optional arguments.
- 3966       3. if *errnum* is nonzero, **error()** shall add the colon and space characters, then  
           3967       the result of **strerror(errnum)**.
- 3968       4. a newline.

3969       If *exitstatus* is nonzero, **error()** shall call **exit(exitstatus)**.

**See Also**

3970       **err()**, **errx()**

**errx****Name**

3971   **errx** — display formatted error message and exit

**Synopsis**

3972   **#include <err.h>**  
 3973   **void errx(int eval, const char \* fmt, ...);**

**Description**

3974   The **errx()** function shall display a formatted error message on the standard error  
 3975   stream. The last component of the program name, a colon character, and a space  
 3976   shall be output. If *fmt* is non-NULL, it shall be used as the format string for the  
 3977   *printf()* family of functions, and the formatted error message, a colon character,  
 3978   and a space shall be output. The output shall be followed by a newline character.

3979   **errx()** does not return, but shall exit with the value of *eval*.

**Return Value**

3980   None.

**Errors**

3981   None.

**See Also**

3982   **error(), err()**

**fcntl****Name**

3983   **fcntl** — file control

**Description**

3984   **fcntl()** is as specified in ISO POSIX (2003), but with differences as listed below.

**Implementation may set O\_LARGEFILE**

3986   According to ISO POSIX (2003), only an application sets **fcntl()** flags, for example  
 3987   **O\_LARGEFILE**. However, this specification also allows an implementation to set the  
 3988   **O\_LARGEFILE** flag in the case where the programming environment is one of  
 3989   **\_POSIX\_V6\_ILP32\_OFFBIG**, **\_POSIX\_V6\_LP64\_OFF64**, **\_POSIX\_V6\_LPBIG\_OFFBIG**.  
 3990   See **getconf** and **c99** in ISO POSIX (2003) for a description of these environments.  
 3991   Thus, calling **fcntl()** with the **F\_GETFL** command may return **O\_LARGEFILE** as well  
 3992   as flags explicitly set by the application in the case that both the implementation and  
 3993   the application support an **off\_t** of at least 64 bits.

## **fflush\_unlocked**

### **Name**

3994      `fflush_unlocked` — non thread safe fflush

### **Description**

3995      `fflush_unlocked()` is the same as `fflush()` except that it need not be thread safe.  
3996      That is, it may only be invoked in the ways which are legal for `getc_unlocked()`.

## **fgetwc\_unlocked**

### **Name**

3997      `fgetwc_unlocked` — non thread safe fgetwc

### **Description**

3998      `fgetwc_unlocked()` is the same as `fgetwc()` except that it need not be thread safe.  
3999      That is, it may only be invoked in the ways which are legal for `getc_unlocked()`.

## flock

### Name

4000      `flock` — apply or remove an advisory lock on an open file

### Synopsis

4001      `int flock(int fd, int operation);`

### Description

4002      `flock()` applies or removes an advisory lock on the open file `fd`. Valid `operation` types are:

4004      `LOCK_SH`

4005      Shared lock. More than one process may hold a shared lock for a given file at a  
4006      given time.

4007      `LOCK_EX`

4008      Exclusive lock. Only one process may hold an exclusive lock for a given file at a  
4009      given time.

4010      `LOCK_UN`

4011      Unlock.

4012      `LOCK_NB`

4013      Don't block when locking. May be specified (by `oring`) along with one of the  
4014      other operations.

4015      A single file may not simultaneously have both shared and exclusive locks.

### Return Value

4016      On success, 0 is returned. On error, -1 is returned and the global variable `errno` is set  
4017      appropriately.

### Errors

4018      `EWOULDBLOCK`

4019      The file is locked and the `LOCK_NB` flag was selected.

**freelocale****Name**

4020      `freelocale` — free a locale object

**Synopsis**

4021      `#include <locale.h>`  
 4022      `void freelocale(locale_t locale);`

**Description**

4023      The `freelocale()` function shall free the locale object `locale`, and release any  
 4024      resources associated with it.

**Return Value**

4025      None.

**Errors**

4026      None defined.

**See Also**

4027      `setlocale()`, `newlocale()`, `duplocale()`, `uselocale()`

**fscanf****Name**

4028      `fscanf` — convert formatted input

**Description**

4029      The `scanf()` family of functions shall behave as described in ISO POSIX (2003),  
 4030      except as noted below.

**Differences**

4031      The `%s`, `%S` and `%[` conversion specifiers shall accept an option length modifier `a`,  
 4032      which shall cause a memory buffer to be allocated to hold the string converted. In  
 4033      such a case, the argument corresponding to the conversion specifier should be a  
 4034      reference to a pointer value that will receive a pointer to the allocated buffer. If there  
 4035      is insufficient memory to allocate a buffer, the function may set `errno` to `ENOMEM`  
 4036      and a conversion error results.

4037      **Note:** This directly conflicts with the ISO C (1999) usage of `%a` as a conversion specifier  
 4038      for hexadecimal float values. While this conversion specifier should be supported, a  
 4039      format specifier such as "`%aseconds`" will have a different meaning on an LSB  
 4040      conforming system.

## fwscanf

### Name

4041 fwscanf — convert formatted input

### Description

4042 The `scanf()` family of functions shall behave as described in ISO POSIX (2003),  
4043 except as noted below.

### Differences

4044 The `%s`, `%S` and `%[` conversion specifiers shall accept an option length modifier `a`,  
4045 which shall cause a memory buffer to be allocated to hold the string converted. In  
4046 such a case, the argument corresponding to the conversion specifier should be a  
4047 reference to a pointer value that will receive a pointer to the allocated buffer. If there  
4048 is insufficient memory to allocate a buffer, the function may set `errno` to `ENOMEM`  
4049 and a conversion error results.

4050 **Note:** This directly conflicts with the ISO C (1999) usage of `%a` as a conversion specifier  
4051 for hexadecimal float values. While this conversion specifier should be supported, a  
4052 format specifier such as "`%aseconds`" will have a different meaning on an LSB  
4053 conforming system.

## getgroupelist

### Name

4054            `getgroupelist — get network group entry`

### Synopsis

```
4055 #include <grp.h>
4056 int getgroupelist(const char * user, gid_t group, gid_t * groups, int *
4057 ngroups);
```

### Description

4058 The `getgroupelist()` function shall fill in the array `groups` with the supplementary  
 4059 groups for the user specified by `user`. On entry, `ngroups` shall refer to an integer  
 4060 containing the maximum number of `gid_t` members in the `groups` array. The group  
 4061 `group` shall also be included. On success, the value referred to by `ngroups` shall be  
 4062 updated to contain the number of `gid_t` objects copied.

### Return Value

4063 On success, if there was sufficient room to copy all the supplementary group  
 4064 identifiers to the array identified by `groups`, `getgroupelist()` shall return the  
 4065 number of `gid_t` objects copied, and the value referenced by `ngroups` shall be  
 4066 updated. If there was not sufficient room to copy all the supplementary group  
 4067 identifiers, `groupelist()` shall return -1, and update the value referenced by  
 4068 `ngroups` to the number actually copied.

4069 If `user` does not refer to a valid user on the system, `getgroupelist()` shall return 0,  
 4070 and set the value referenced by `ngroups` to 0.

### Errors

4071 None defined.

### See Also

4072            `getgroups( )`

## getloadavg

### Name

4073            `getloadavg — get system load averages`

### Synopsis

```
4074 #include <stdlib.h>
4075 int getloadavg(double loadavg[], int nelem);
```

### Description

4076 `getloadavg( )` returns the number of processes in the system run queue averaged  
 4077 over various periods of time. Up to `nelem` samples are retrieved and assigned to  
 4078 successive elements of `loadavg[]`. The system imposes a maximum of 3 samples,  
 4079 representing averages over the last 1, 5, and 15 minutes, respectively.

## getopt

### Name

4080 getopt — parse command line options

### Synopsis

```
4081 #include <unistd.h>
4082 int getopt(int argc, char * const argv[], const char * optstring);
4083 extern char *optarg;
```

4084       extern int optind, opterr, getopt;

## Description

4085       The `getopt()` function shall parse command line arguments as described in ISO  
 4086       POSIX (2003), with the following exceptions, where LSB and POSIX specifications  
 4087       vary. LSB systems shall implement the modified behaviors described below.

### Argument Ordering

4089       The `getopt()` function can process command line arguments referenced by `argv` in  
 4090       one of three ways:

4091       PERMUTE

4092       the order of arguments in `argv` is altered so that all options (and their  
 4093       arguments) are moved in front of all of the operands. This is the default  
 4094       behavior.

4095       **Note:** This behavior has undefined results if `argv` is not modifiable. This is to support  
 4096       historic behavior predating the use of `const` and ISO C (1999). The function  
 4097       prototype was aligned with ISO POSIX (2003) despite the fact that it modifies `argv`,  
 4098       and the library maintainers are unwilling to change this.

4099       REQUIRE\_ORDER

4100       The arguments in `argv` are processed in exactly the order given, and option  
 4101       processing stops when the first non-option argument is reached, or when the  
 4102       element of `argv` is "--". This ordering can be enforced either by setting the  
 4103       environment variable `POSIXLY_CORRECT`, or by setting the first character of  
 4104       `optstring` to '+'.

4105       RETURN\_IN\_ORDER

4106       The order of arguments is not altered, and all arguments are processed.  
 4107       Non-option arguments (operands) are handled as if they were the argument to  
 4108       an option with the value 1 ('\001'). This ordering is selected by setting the first  
 4109       character of `optstring` to '-'.

### Option Characteristics

4111       LSB specifies that:

- 4112       • an element of `argv` that starts with "-" (and is not exactly "-" or "--") is an option  
     4113       element.
- 4114       • characters of an option element, aside from the initial "-", are option characters.

4115       POSIX specifies that:

- 4116       • applications using `getopt()` shall obey the following syntax guidelines:
  - 4117           • option name is a single alphanumeric character from the portable character set
  - 4118           • option is preceded by the '-' delimiter character
  - 4119           • options without option-arguments should be accepted when grouped behind  
         4120            one '-' delimiter
  - 4121           • each option and option-argument is a separate argument
  - 4122           • option-arguments are not optional
  - 4123           • all options should precede operands on the command line

- 4124     • the argument "--" is accepted as a delimiter indicating the end of options and the  
 4125        consideration of subsequent arguments, if any, as operands
- 4126     • historical implementations of `getopt()` support other characters as options as an  
 4127        allowed extension, but applications that use extensions are not maximally  
 4128        portable.
- 4129     • support for multi-byte option characters is only possible when such characters can  
 4130        be represented as type `int`.
- 4131     • applications that call any utility with a first operand starting with '-' should  
 4132        usually specify "--" to mark the end of the options. Standard utilities that do not  
 4133        support this guideline indicate that fact in the OPTIONS section of the utility  
 4134        description.

4135     **Extensions**

4136     *LSB* specifies that:

- 4137        • if a character is followed by two colons, the option takes an optional argument; if  
 4138        there is text in the current `argv` element, it is returned in `optarg`, otherwise  
 4139        `optarg` is set to 0.
- 4140        • if `optstring` contains `w` followed by a semi-colon (`:`), then `-w foo` is treated as the  
 4141        long option `--foo`.

4142     **Note:** See `getopt_long()` for a description of long options.

- 4143        • The first character of `optstring` shall modify the behavior of `getopt()` as follows:
  - 4144            • if the first character is '+', then `REQUIRE_ORDER` processing shall be in effect (see  
 4145            above)
  - 4146            • if the first character is '!', then `RETURN_IN_ORDER` processing shall be in effect  
 4147            (see above)
  - 4148            • if the first character is ':', then `getopt()` shall return ':' instead of '?' to indicate a  
 4149            missing option argument, and shall not print any diagnostic message to `stderr`.

4150     *POSIX* specifies that:

- 4151        • the `-w` option is reserved for implementation extensions.

4152     **Return Values**

4153     *LSB* specifies the following additional `getopt()` return values:

- 4154        • '\001' is returned if `RETURN_IN_ORDER` argument ordering is in effect, and the next  
 4155        argument is an operand, not an option. The argument is available in `optarg`.

4156     Any other return value has the same meaning as for *POSIX*.

4157     *POSIX* specifies the following `getopt()` return values:

- 4158        • the next option character is returned, if found successfully.
- 4159        • ':' is returned if a parameter is missing for one of the options and the first character  
 4160        of `optstring` is ':'.
- 4161        • '?' is returned if an unknown option character not in `optstring` is encountered, or  
 4162        if `getopt()` detects a missing argument and the first character of `optstring` is not  
 4163        ':'.
- 4164        • -1 is returned for the end of the option list.

4165       **Environment Variables**

4166       LSB specifies that:

- if the variable `POSIXLY_CORRECT` is set, option processing stops as soon as a non-option argument is encountered.
- the variable `_[_PID]_GNU_nonoption_argv_flags_` (where `[PID]` is the process ID for the current process), contains a space separated list of arguments that should not be treated as arguments even though they appear to be so.

4172       **Rationale:** This was used by bash 2.0 to communicate to *GNU libc* which arguments  
 4173       resulted from wildcard expansion and so should not be considered as options. This  
 4174       behavior was removed in bash version 2.01, but the support remains in *GNU libc*.

4175       This behavior is DEPRECATED in this version of the LSB; future revisions of this  
 4176       specification may not include this requirement.

**getopt\_long****Name**4177       `getopt_long` — parse command line options**Synopsis**

```
4178 #define _GNU_SOURCE
4179 #include <getopt.h>
4180 int getopt_long(int argc, char * const argv[], const char * opstring, const
4181 struct option * longopts, int * longindex);
```

**Description**

4182       `getopt_long()` works like `getopt()` except that it also accepts long options, started  
 4183       out by two dashes. Long option names may be abbreviated if the abbreviation is  
 4184       unique or is an exact match for some defined option. A long option may take a  
 4185       parameter, of the form `--arg=param` or `--arg param`.

4186       `longopts` is a pointer to the first element of an array of `struct option` declared in  
 4187       `getopt.h` as:

```
4188 struct option {
4189     const char *name;
4190     int has_arg;
4191     int *flag;
4192     int val;
4193 };
```

4194       The fields in this structure have the following meaning:

4195       `name`

4196       The name of the long option.

4197       `has_arg`

4198       One of:

4199                   argument (or 0) if the option does not take an argument,  
 uired\_argument (or 1) if the option requires an argument, or  
 ional\_argument (or 2) if the option takes an optional argument.

4200                   *flag*  
 4201                   specifies how results are returned for a long option. If *flag* is NULL, then  
 4202                   getopt\_long() shall return *val*. (For example, the calling program may set *val*  
 4203                   to the equivalent short option character.) Otherwise, getopt\_long() returns 0,  
 4204                   and *flag* shall point to a variable which shall be set to *val* if the option is found,  
 4205                   but left unchanged if the option is not found.

4206                   *val*  
 4207                   The value to return, or to load into the variable pointed to by *flag*.

## Return Value

4208                   getopt\_long() returns the option character if a short option was found successfully,  
 4209                   or ":" if there was a missing parameter for one of the options, or "?" for an unknown  
 4210                   option character, or -1 for the end of the option list.  
 4211                   For a long option, getopt\_long() returns *val* if *flag* is NULL, and 0 otherwise. Error  
 4212                   and -1 returns are the same as for getopt(), plus "?" for an ambiguous match or an  
 4213                   extraneous parameter.

## getopt\_long\_only

### Name

4214                   **getopt\_long\_only** — parse command line options

### Synopsis

4215                   #define \_GNU\_SOURCE

```
4216     #include <getopt.h>
4217     int getopt_long_only(int argc, char * const argv[], const char * optstring,
4218     const struct option * longopts, int * longindex);
```

## Description

4219     `getopt_long_only()` is like `getopt_long()`, but `"-"` as well as `--` can indicate a  
 4220     long option. If an option that starts with `"-"` (not `--`) doesn't match a long option, but  
 4221     does match a short option, it is parsed as a short option instead.

4222     **Note:** The `getopt_long_only()` function is intended only for supporting certain  
 4223     programs whose command line syntax was designed before the Utility Syntax  
 4224     Guidelines of ISO POSIX (2003) were developed. New programs should generally call  
 4225     `getopt_long()` instead, which provides the `--option` syntax for long options, which is  
 4226     preferred by GNU and consistent with ISO POSIX (2003).

## Return Value

4227     `getopt_long_only()` returns the option character if the option was found  
 4228     successfully, or `:` if there was a missing parameter for one of the options, or `?"` for  
 4229     an unknown option character, or `-1` for the end of the option list.

4230     `getopt_long_only()` also returns the option character when a short option is  
 4231     recognized. For a long option, they return `val` if `flag` is `NULL`, and `0` otherwise. Error  
 4232     and `-1` returns are the same as for `getopt()`, plus `?"` for an ambiguous match or an  
 4233     extraneous parameter.

## getsockopt

### Name

4234     `getsockopt` — get socket options

### Synopsis

```
4235     #include <sys/socket.h>
```

```
4236 #include <netinet/ip.h>
4237 int getsockopt(int socket, int level, int option_name, void * restrict
4238 option_value, socklen_t * restrict option_len);
```

## Description

4239 The `getsockopt()` function shall behave as specified in ISO POSIX (2003), with the  
 4240 following extensions.

### IP Protocol Level Options

4242 If the `level` parameter is  `IPPROTO_IP`, the following values shall be supported for  
 4243 `option_name` (see RFC 791:Internet Protocol for further details):

4244 `IP_OPTIONS`

4245 Get the Internet Protocol options sent with every packet from this socket. The  
 4246 `option_value` shall point to a memory buffer in which the options shall be  
 4247 placed; on entry `option_len` shall point to an integer value indicating the  
 4248 maximum size of the memory buffer, in bytes. On successful return, the value  
 4249 referenced by `option_len` shall be updated to the size of data copied to the  
 4250 buffer. For IPv4, the maximum length of options is 40 bytes.

4251 `IP_TTL`

4252 Get the current unicast Internet Protocol Time To Live value used when sending  
 4253 packets with this socket. The `option_value` shall point to a buffer large enough  
 4254 to hold the time to live value (at least 1 byte), and `option_len` shall point to an  
 4255 integer value holding the maximum size of that buffer. On successful return, the  
 4256 value referenced by `option_len` shall be updated to contain the number of  
 4257 bytes copied into the buffer, which shall be no larger than the initial value, and  
 4258 `option_value` shall point to an integer containing the time to live value.

4259 `IP_TOS`

4260 Get the Internet Protocol type of service indicator used when sending packets  
 4261 with this socket. The `option_value` shall point to a buffer large enough to hold  
 4262 the type of service indicator (at least 1 byte), and `option_len` shall point to an  
 4263 integer value holding the maximum size of that buffer. On successful return, the  
 4264 value referenced by `option_len` shall be updated to contain the number of  
 4265 bytes copied into the buffer, which shall be no larger than the initial value, and  
 4266 `option_value` shall point to an integer containing the time to live value.

## gettext

### Name

4267    `gettext — search message catalogs for a string`

### Synopsis

4268    `#include <libintl.h>`  
4269    `char * gettext(const char * msgid);`

### Description

4270    The `gettext()` function shall search the currently selected message catalogs for a  
4271    string identified by the string `msgid`. If a string is located, that string shall be  
4272    returned.

4273    The `gettext()` function is equivalent to `dcgettext(NULL, msgid, LC_MESSAGES)`.

### Return Value

4274    If a string is found in the currently selected message catalogs for `msgid`, then a  
4275    pointer to that string shall be returned. Otherwise, a pointer to `msgid` shall be  
4276    returned.

4277    Applications shall not modify the string returned by `gettext()`.

### Errors

4278    None.

4279    The `gettext()` function shall not modify `errno`.

### See Also

4280    `dgettext`, `ngettext`, `dngettext`, `dcgettext`, `dcngettext`, `textdomain`, `bindtextdomain`,  
4281    `bind_textdomain_codeset`

## getutent

### Name

4282 getutent — access user accounting database entries

### Synopsis

4283 #include <utmp.h>  
4284 struct utmp \*getutent(void);

### Description

4285 The `getutent()` function shall read the next entry from the user accounting  
4286 database.

### Return Value

4287 Upon successful completion, `getutent()` shall return a pointer to a `utmp` structure  
4288 containing a copy of the requested entry in the user accounting database. Otherwise,  
4289 a null pointer shall be returned. The return value may point to a static area which is  
4290 overwritten by a subsequent call to `getutent()`.

### Errors

4291 None defined.

## getutent\_r

### Name

4292 `getutent_r` — access user accounting database entries

### Synopsis

4293 `int getutent_r(struct utmp * buffer, struct utmp ** result);`

### Description

4294 The `getutent_r()` function is a reentrant version of the `getutent()` function. On  
4295 entry, `buffer` should point to a user supplied buffer to which the next entry in the  
4296 database will be copied, and `result` should point to a location where the result will  
4297 be stored.

### Return Value

4298 On success, `getutent_r()` shall return 0 and set the location referenced by `result`  
4299 to a pointer to `buffer`. Otherwise, `getutent_r()` shall return -1 and set the location  
4300 referenced by `result` to NULL.

## glob64

### Name

4301      glob64 — find pathnames matching a pattern (Large File Support)

### Synopsis

```
4302 #include <glob.h>
4303 int glob64(const char * pattern, int flags, int (*errfunc) (const char *, int),
4304   glob64_t * pglob);
```

### Description

4305 The `glob64()` function is a large-file version of the `glob()` defined in ISO POSIX  
 4306 (2003). It shall search for pathnames matching *pattern* according to the rules used  
 4307 by the shell, `/bin/sh`. No tilde expansion or parameter substitution is done; see  
 4308 `wordexp()`.

4309 The results of a `glob64()` call are stored in the structure pointed to by *pglob*, which  
 4310 is a `glob64_t` declared in `glob.h` with the following members:

```
4311 typedef struct
4312 {
4313     size_t gl_pathc;
4314     char **gl_pathv;
4315     size_t gl_offs;
4316     int gl_flags;
4317     void (*gl_closedir) (void *);
4318     struct dirent64 *(*gl_readdir64) (void *);
4319     void *(*gl_opendir) (const char *);
4320     int (*gl_lstat) (const char *, struct stat *);
4321     int (*gl_stat) (const char *, struct stat *);
4322 }
```

```

4323     glob64_t;

4324     Structure members with the same name as corresponding members of a glob_t as
4325     defined in ISO POSIX (2003) shall have the same purpose.

4326     Other members are defined as follows:

4327     gl_flags
4328         reserved for internal use

4329     gl_closedir
4330         pointer to a function capable of closing a directory opened by gl_opendir

4331     gl_readdir64
4332         pointer to a function capable of reading entries in a large directory

4333     gl_opendir
4334         pointer to a function capable of opening a large directory

4335     gl_stat
4336         pointer to a function capable of returning file status for a large file

4337     gl_lstat
4338         pointer to a function capable of returning file status information for a large file
4339         or symbolic link

4340     A large file or large directory is one with a size which cannot be represented by a
4341     variable of type off_t.

```

## Return Value

```

4342     On success, 0 is returned. Other possible returns are:

4343     GLOB_NOSPACE
4344         out of memory

4345     GLOB_ABORTED
4346         read error

4347     GLOB_NOMATCH
4348         no match found

```

**globfree64****Name**

4349      globfree64 — free memory from glob64() (Large File Support)

**Synopsis**

4350      #include <glob.h>  
 4351      void globfree64(glob64\_t \* pglob);

**Description**

4352      globfree64() frees the dynamically allocated storage from an earlier call to  
 4353      glob64().

4354      globfree64() is a 64-bit version of globfree().

**initgroups****Name**

4355      initgroups — initialize the supplementary group access list

**Synopsis**

4356      #include <grp.h>  
 4357      #include <sys/types.h>  
 4358      int initgroups(const char \* user, gid\_t group);

**Description**

4359      If the process has appropriate privilege, the initgroups() function shall initialize  
 4360      the Supplementary Group IDs for the current process by reading the group database  
 4361      and using all groups of which *user* is a member. The additional group *group* is also  
 4362      added to the list.

**Return Value**

4363      On success, 0 is returned. On error, -1 is returned and the global variable errno is set  
 4364      appropriately.

**Errors**

4365      EPERM

4366      The calling process does not have sufficient privileges.

4367      ENOMEM

4368      Insufficient memory to allocate group information structure.

**See Also**

4369      setgroups()

## ioctl

### Name

4370    `ioctl` — control device

### Synopsis

4371    `#include <sys/iotcl.h>`  
 4372    `int ioctl (int fildes , int request , ...);`

### Description

4373    The `ioctl()` function shall manipulate the underlying device parameters of special  
 4374    files. *fildes* shall be an open file descriptor referring to a special file. The `ioctl()`  
 4375    function shall take three parameters; the type and value of the third parameter is  
 4376    dependent on the device and *request*.

4377    Conforming LSB applications shall not call `ioctl()` except in situations explicitly  
 4378    stated in this specification.

### Return Value

4379    On success, 0 is returned. An `ioctl()` may use the return value as an output  
 4380    parameter and return a non-negative value on success. On error, -1 is returned and  
 4381    the global variable `errno` is set appropriately.

### Errors

4382    EBADF

4383       *fildes* is not a valid descriptor.

4384    EFAULT

4385       The third parameter references an inaccessible memory area.

4386    ENOTTY

4387       *fildes* is not associated with a character special device.

4388    ENOTTY

4389       The specified request does not apply to the kind of object that *fildes*  
 4390       references.

4391    EINVAL

4392       *request* or the third parameter is not valid.

### Relationship to POSIX (Informative)

4393    It should be noted that ISO POSIX (2003) contains an interface named `ioct1()`. The  
 4394    LSB only defines behavior when *fildes* refers to a socket (see `sockio`) or terminal  
 4395    device (see `ttyio`), while ISO POSIX (2003) only defines behavior when *fildes* refers  
 4396    to a STREAMS device. An implementation may support both behaviors; the LSB  
 4397    does not require any STREAMS support.

## sockio

### Name

4398        sockio — socket ioctl commands

### Synopsis

4399        #include <sys/ioctl.h>  
4400        #include <sys/socket.h>  
4401        #include <net/if.h>

```
4402 #include <netinet/in.h>
4403 int ioctl(int sockfd, int request, void * argp);
```

## Description

4404  
4405  
4406 Socket `ioctl()` commands are a subset of the `ioctl()` calls, which can perform a variety of functions on sockets. `sockfd` shall be an open file descriptor referring to a socket (see the `socket()` or `accept()` functions).

4407  
4408 Socket `ioctl()` commands apply to the underlying network interfaces, and affect the entire system, not just the file descriptor used to issue the `ioctl()`.

4409 The following values for `request` are accepted:

4410 SIOCGIFCONF (Deprecated)

4411 Get the interface configuration list for the system.

4412 **Note:** The `SIOCGIFCONF` interface is superceded by the `if_nameindex()` family of  
4413 functions (see ISO POSIX (2003)). A future version of this specification may  
4414 withdraw this value for `request`.

4415 `argp` shall point to a `ifconf` structure, as described in `<net/if.h>`. Before  
4416 calling, the caller shall set the `ifc_ifcu.ifcu_req` field to point to an array of  
4417 `ifreq` structures, and set `ifc_len` to the size in bytes of this allocated array.  
4418 Upon return, `ifc_len` will contain the size in bytes of the array which was  
4419 actually used. If it is the same as the length upon calling, the caller should  
4420 assume that the array was too small and try again with a larger array.

4421 On success, `SIOCGIFCONF` shall return a nonnegative value.

4422 **Rationale:** Historical UNIX systems disagree on the meaning of the return value.

4423 SIOCGIFFLAGS

4424 Get the interface flags for the indicated interface. `argp` shall point to a `ifreq`  
4425 structure. Before calling, the caller should fill in the `ifr_name` field with the  
4426 interface name, and upon return, the `ifr_ifru.ifru_flags` field is set with the  
4427 interface flags.

4428 SIOCGIFADDR

4429 Get the interface address for the given interface. `argp` shall point to a `ifreq`  
4430 structure. Before calling, the caller should fill in the `ifr_name` field with the  
4431 interface name, and upon return, the `ifr_ifru.ifru_addr` field is set with the  
4432 interface address.

4433 SIOCGIFBRDADDR

4434 Get the interface broadcast address for the given interface. `argp` shall point to a  
4435 `ifreq` structure. Before calling, the caller should fill in the `ifr_name` field with the  
4436 interface name, and upon return, the `ifr_ifru.ifru_broadcast` field is set with the  
4437 interface broadcast address.

4438 SIOCGIFNETMASK

4439 Get the network mask for the given interface. `argp` shall point to a `ifreq`  
4440 structure. Before calling, the caller should fill in the `ifr_name` field with the  
4441 interface name, and upon return, the `ifr_ifru.ifru_netmask` field is set with the  
4442 network mask.

4443 SIOCGIFMTU  
 4444       Get the Maximum Transmission Unit (MTU) size for the given interface. *argp*  
 4445       shall point to a *ifreq* structure. Before calling, the caller should fill in the  
 4446       *ifr\_name* field with the interface name, and upon return, the  
 4447       *ifr\_ifru.ifru\_mtu* field is set with the MTU.

4448 FIONREAD  
 4449       Get the amount of queued unread data in the receive buffer. *argp* shall point to  
 4450       an integer where the result is to be placed.

4451       **Note:** Some implementations may also support the use of FIONREAD on other types of file  
 4452       descriptor. However, the LSB only specifies its behavior for a socket related file  
 4453       descriptor.

## Return Value

4454 On success, if *request* is SIOCGIFCONF, a non-negative integer shall be returned. If  
 4455 request is not SIOCGIFCONF, on success 0 is returned. On error, -1 is returned and  
 4456 the global variable *errno* is set appropriately.

## Errors

4457 EBADF  
 4458       *sockfd* is not a valid descriptor.

4459 EFAULT  
 4460       *argp* references an inaccessible memory area.

4461 ENOTTY  
 4462       The specified *request* does not apply to the kind of object that the descriptor  
 4463       *sockfd* references.

4464 EINVAL  
 4465       Either *request* or *argp* is invalid.

4466 ENOTCONN  
 4467       The operation is only defined on a connected socket, but the socket wasn't  
 4468       connected.

## ttyio

### Name

4469 ttyio — tty ioctl commands

### Synopsis

4470 #include <sys/ioctl.h>

```
4471 #include <fcntl.h>
4472 int ioctl(int fd, unsigned long request, int * argp);
```

## Description

4473 Tty *ioctl* commands are a subset of the *ioctl()* calls, which can perform a variety of  
 4474 functions on tty devices. *fd* shall be an open file descriptor referring to a terminal  
 4475 device.

4476 The following *ioctl()*s are provided:

4477 TIOCGWINSZ

4478 Get the size attributes of the terminal or pseudo-terminal identified by *fd*. On  
 4479 entry, *argp* shall reference a *winsize* structure. On return, the structure will  
 4480 have *ws\_row* set to the number of rows of text (i.e. lines of text) that can be  
 4481 viewed on the device, and *ws\_col* set to the number of columns (i.e. text width).

4482 **Note:** The number of columns stored in *ws\_col* assumes that the terminal device is using  
 4483 a mono-spaced font.

## Return Value

4484 On success, 0 is returned. On error, -1 is returned and the global variable *errno* is set  
 4485 appropriately.

## Errors

4486 EBADF

4487 *fd* is not a valid descriptor.

4488 EFAULT

4489 *argp* references an inaccessible memory area.

4490 EINVAL

4491 *request* and *argp* are not valid.

**kill****Name**

4492       **kill** — send a signal

**Synopsis**

4493       

```
#include <signal.h>
4494       int kill(pid_t pid, int sig);
```

**Description**

4495       **kill()** is as specified in the *ISO POSIX (2003)*, but with differences as listed below.

**Process ID -1 doesn't affect calling process**

4497       If *pid* is specified as *-1*, *sig* shall not be sent to the calling process. Other than this,  
 4498       the rules in the *ISO POSIX (2003)* apply.

4499       **Rationale:** This was a deliberate Linus decision after an unpopular experiment in  
 4500       including the calling process in the 2.5.1 kernel. See "What does it mean to signal  
 4501       everybody?", Linux Weekly News, 20 December 2001,  
 4502       <http://lwn.net/2001/1220/kernel.php3>

**link****Name**

4503       **link** — create a link to a file

**Synopsis**

4504       

```
#include <unistd.h>
4505       int link(const char * path1, const char * path2);
```

**Description**

4506       The **link()** function shall behave as specified in *ISO POSIX (2003)*, except with  
 4507       differences as listed below.

**Need Not Follow Symlinks**

4509       ISO POSIX (2003) specifies that pathname resolution shall follow symbolic links  
 4510       during pathname resolution unless the function is required to act on the symbolic  
 4511       link itself, or certain arguments direct that the function act on the symbolic link itself.  
 4512       The **link()** function in ISO POSIX (2003) contains no such requirement to operate  
 4513       on a symbolic link. However, a conforming LSB implementation need not follow a  
 4514       symbolic link for the *path1* argument.

## mbsnrtowcs

### Name

4515 mbsnrtowcs — convert a multibyte string to a wide character string

### Synopsis

```
4516 #include <wchar.h>
4517 size_t mbsnrtowcs(wchar_t * dest, const char * * src, size_t nms, size_t len,
4518 mbstate_t * ps);
```

### Description

4519 mbsnrtowcs() is like mbsrtowcs(), except that the number of bytes to be converted,  
 4520 starting at *src*, is limited to *nms*.

4521 If *dest* is not a NULL pointer, mbsnrtowcs() converts at most *nms* bytes from the  
 4522 multibyte string *src* to a wide-character string starting at *dest*. At most, *len* wide  
 4523 characters are written to *dest*. The state *ps* is updated.

4524 The conversion is effectively performed by repeatedly calling:

4525

4526            `mbrtowc(dest, *src, n, ps)`

4527            where *n* is some positive number, as long as this call succeeds, and then  
 4528            incrementing *dest* by one and *src* by the number of bytes consumed.

4529            The conversion can stop for three reasons:

- 4530            • An invalid multibyte sequence has been encountered. In this case *src* is left  
 4531            pointing to the invalid multibyte sequence, `(size_t)(-1)` is returned, and `errno` is  
 4532            set to `EILSEQ`.
- 4533            • The *nms* limit forces a stop, or *len* non-L'\0' wide characters have been stored at  
 4534            *dest*. In this case, *src* is left pointing to the next multibyte sequence to be  
 4535            converted, and the number of wide characters written to *dest* is returned.
- 4536            • The multibyte string has been completely converted, including the terminating  
 4537            '\0' (which has the side effect of bringing back *ps* to the initial state). In this case,  
 4538            *src* is set to `NULL`, and the number of wide characters written to *dest*, excluding  
 4539            the terminating L'\0' character, is returned.

4540            If *dest* is `NULL`, *len* is ignored, and the conversion proceeds as above, except that the  
 4541            converted wide characters are not written out to memory, and that no destination  
 4542            length limit exists.

4543            In both of the above cases, if *ps* is a `NULL` pointer, a static anonymous state only  
 4544            known to `mbsnrtowcs()` is used instead.

4545            The programmer shall ensure that there is room for at least *len* wide characters at  
 4546            *dest*.

## Return Value

4547            `mbsnrtowcs()` returns the number of wide characters that make up the converted  
 4548            part of the wide character string, not including the terminating null wide character.  
 4549            If an invalid multibyte sequence was encountered, `(size_t)(-1)` is returned, and the  
 4550            global variable `errno` is set to `EILSEQ`.

## Notes

4551            The behavior of `mbsnrtowcs()` depends on the `LC_CTYPE` category of the current  
 4552            locale.

4553            Passing `NULL` as *ps* is not multi-thread safe.

## **memmem**

### Name

4554            `memmem — locate bytes`

### Synopsis

4555            `#define _GNU_SOURCE`

```
4556 #include <string.h>
4557 void * memmem(const void * haystack, size_t haystacklen, const void * needle,
4558 size_t needlelen);
```

## Description

4559 memmem( ) finds the start of the first occurrence of the byte array referenced by  
 4560 *needle* of length *needlelen* in the memory area *haystack* of length *haystacklen*.

## Return Value

4561 memmem( ) returns a pointer to the beginning of the byte array, or NULL if the byte  
 4562 array is not found.

## Notes

4563 Earlier versions of the C library (prior to glibc 2.1) contained a memmem( ) with  
 4564 various problems, and application developers should treat this function with care.

## memrchr

### Name

4565 memrchr — scan memory for a character

### Synopsis

```
4566 #include <string.h>
4567 void * memrchr(const void * s, int c, size_t n);
```

## Description

4568 The memrchr( ) function shall locate the last occurrence of *c* (converted to an  
 4569 unsigned char) in the initial *n* bytes (each interpreted as an unsigned char) of the  
 4570 object pointed to by *s*.

## Return Value

4571 The memrchr( ) shall return a pointer to the located byte, or a null pointer if the byte  
 4572 does not occur in the object.

## Errors

4573 No errors are defined.

## See Also

4574 memchr( )

## newlocale

### Name

4575 newlocale — allocate a locale object

### Synopsis

4576 

```
#include <locale.h>
4577 locale_t newlocale(int category_mask, const char * locale, locale_t base);
```

### Description

4578 The `newlocale()` function shall initialize a locale object. If `base` is `NULL`, then  
 4579 `newlocale()` shall first allocate the object; otherwise it shall use the locale object  
 4580 referenced by `base`.

4581 The object shall be initialized for the locale named by `locale`, and for the categories  
 4582 selected in `category_mask`. The `category_mask` value is a bitwise inclusive OR of  
 4583 the required `LC_name_MASK` values, or the value `LC_ALL_MASK`.

### Return Value

4584 On success, the `newlocale()` function shall return the initialized locale object.  
 4585 Otherwise, it shall return `NULL`, and set `errno` to indicate the error.

### Errors

4586 The `newlocale()` function shall fail if:

4587 `ENOMEM`

4588       Insufficient memory.

4589 `EINVAL`

4590       An invalid `category_mask` was provided, or the `locale` was `NULL`.

### Application Usage (Informative)

4591 The only portable way to allocate a locale object is to call `newlocale()` with a `NULL`  
 4592 `base`. The allocated object may be reinitialized to a new locale by passing it back to  
 4593 `newlocale()`. The new object may be released by calling `freelocale()`.

### See Also

4594 `setlocale()`, `freelocale()`, `duplocale()`, `uselocale()`

## ngettext

### Name

4595           ngettext — search message catalogs for plural string

### Synopsis

```
4596 #include <libintl.h>
4597 char * ngettext(const char * msgid1, const char * msgid2, unsigned long int
4598 n);
```

### Description

4599           The `ngettext()` function shall search the currently selected message catalogs for a  
 4600           string matching the singular string `msgid1`. If a string is located, and if `n` is 1, that  
 4601           string shall be returned. If `n` is not 1, a pluralized version (dependent on `n`) of the  
 4602           string shall be returned.

4603           The `ngettext()` function is equivalent to `dcngettext(NULL, msgid1, msgid2, n,`  
 4604           `LC_MESSAGES)()`.

### Return Value

4605           If a string is found in the currently selected message catalogs for `msgid1`, then if `n` is  
 4606           1 a pointer to the located string shall be returned. If `n` is not 1, a pointer to an  
 4607           appropriately pluralized version of the string shall be returned. If no message could  
 4608           be found in the currently selected mesage catalogs, then if `n` is 1, a pointer to `msgid1`  
 4609           shall be returned, otherwise a pointer to `msgid2` shall be returned.

4610           Applications shall not modify the string returned by `ngettext()`.

### Errors

4611           None.

4612           The `ngettext()` function shall not modify `errno`.

### See Also

4613           `gettext`, `dgettext`, `ngettext`, `dngettext`, `dcgettext`, `dcngettext`, `textdomain`,  
 4614           `bindtextdomain`, `bind_textdomain_codeset`

## pmap\_getport

### Name

4615 pmap\_getport — find the port number assigned to a service registered with a  
 4616 portmapper.

### Synopsis

```
4617 #include <rpc/pmap_clnt.h>
4618 u_short * pmap_getport(struct sockaddr_in * address, const u_long program,
4619 const u_long * version, u_int protocol);
```

### Description

4620 The `pmap_getport()` function shall return the port number assigned to a service  
 4621 registered with a RPC Binding service running on a given target system, using the  
 4622 protocol described in RFC 1833: Binding Protocols for ONC RPC Version 2. The  
 4623 `pmap_getport()` function shall be called given the RPC program number *program*,  
 4624 the program version *version*, and transport protocol *protocol*. Conforming  
 4625 implementations shall support both `IPPROTO_UDP` and `IPPROTO_TCP` protocols. On  
 4626 entry, *address* shall specify the address of the system on which the portmapper to  
 4627 be contacted resides. The value of *address->sin\_port* shall be ignored, and the  
 4628 standard value for the portmapper port shall always be used.

4629 **Note:** Security and network restrictions may prevent a conforming application from  
 4630 contacting a remote RPC Binding Service.

### Return Value

4631 On success, the `pmap_getport()` function shall return the port number in host byte  
 4632 order of the RPC application registered with the remote portmapper. On failure, if  
 4633 either the program was not registered or the remote portmapper service could not be  
 4634 reached, the `pmap_getport()` function shall return 0. If the remote portmap service  
 4635 could not be reached, the status is left in the global variable `rpc_createerr`.

## pmap\_set

### Name

4636 `pmap_set` — establishes mapping to machine's RPC Bind service.

### Synopsis

```
4637 #include <rpc/pmap_clnt.h>
4638 bool_t pmap_set(const u_long program, const u_long version, int protocol,
4639 u_short port);
```

### Description

4640 `pmap_set()` establishes a mapping between the triple  
 4641 `[program,version,protocol]` and *port* on the machine's RPC Bind service. The  
 4642 value of *protocol* is most likely `IPPROTO_UDP` or `IPPROTO_TCP`. Automatically done  
 4643 by `svc_register()`.

### Return Value

4644 `pmap_set()` returns non-zero if it succeeds, 0 otherwise.

## pmap\_unset

### Name

4645 pmap\_unset — destroys RPC Binding

### Synopsis

```
4646 #include <rpc/pmap_clnt.h>
4647
4648 bool_t pmap_unset(u_long program, u_long versnum);
```

### Description

4650 As a user interface to the RPC Bind service, `pmap_unset()` destroys all mapping  
 4651 between the triple [*programnum,versnum,\**] and ports on the machine's RPC Bind  
 4652 service.

### Return Value

4653 `pmap_unset()` returns non-zero if it succeeds, zero otherwise.

## psignal

### Name

4654 `psignal` — print signal message

### Synopsis

```
4655 #include <signal.h>
4656 void psignal(int sig, const char * s);
4657 extern const char *const sys_siglist[]
```

### Description

4658 The `psignal()` function shall display a message on the `stderr` stream. If *s* is not the  
 4659 null pointer, and does not point to an empty string (e.g. "\0"), the message shall  
 4660 consist of the string *s*, a colon, a space, and a string describing the signal number  
 4661 *sig*; otherwise `psignal()` shall display only a message describing the signal  
 4662 number *sig*. If *sig* is invalid, the message displayed shall indicate an unknown  
 4663 signal.

4664 The array `sys_siglist` holds the signal description strings indexed by signal  
 4665 number.

### Return Value

4666 `psignal()` returns no value.

**regexec****Name**

4667      `regexec` — regular expression matching

**Description**

4668      The `regexec()` function shall behave as specified in *ISO POSIX (2003)*, except with  
4669      differences as listed below.

**Differences**

4670      Certain aspects of regular expression matching are optional; see Internationalization  
4671      and Regular Expressions.

**scanf****Name**

4673      `scanf` — convert formatted input

**Description**

4674      The `scanf()` family of functions shall behave as described in *ISO POSIX (2003)*,  
4675      except as noted below.

**Differences**

4676      The `%s`, `%S` and `%[` conversion specifiers shall accept an option length modifier `a`,  
4677      which shall cause a memory buffer to be allocated to hold the string converted. In  
4678      such a case, the argument corresponding to the conversion specifier should be a  
4679      reference to a pointer value that will receive a pointer to the allocated buffer. If there  
4680      is insufficient memory to allocate a buffer, the function may set `errno` to `ENOMEM`  
4681      and a conversion error results.

4682      **Note:** This directly conflicts with the ISO C (1999) usage of `%a` as a conversion specifier  
4683      for hexadecimal float values. While this conversion specifier should be supported, a  
4684      format specifier such as "`%aseconds`" will have a different meaning on an LSB  
4685      conforming system.

**setbuffer****Name**

4686      `setbuffer` — stream buffering operation

**Synopsis**

```
4687 #include <stdio.h>
4688 void setbuffer(FILE * stream, char * buf, size_t size);
```

**Description**

4689      `setbuffer()` is an alias for the call to `setvbuf()`. It works the same, except that the  
4690      size of the buffer in `setbuffer()` is up to the caller, rather than being determined by  
4691      the default `BUFSIZ`.

**setgroups****Name**

4692       **setgroups** — set list of supplementary group IDs

**Synopsis**

```
4693 #include <grp.h>
4694 int setgroups(size_t size, const gid_t * list);
```

**Description**

4695       If the process has appropriate privilege, the **setgroups()** function shall set the  
 4696       supplementary group IDs for the current process. *list* shall reference an array of  
 4697       *size* group IDs. A process may have at most **NGROUPS\_MAX** supplementary group  
 4698       IDs.

**Return Value**

4699       On successful completion, 0 is returned. On error, -1 is returned and the **errno** is set  
 4700       to indicate the error.

**Errors**

4701       **EFAULT**

4702       *list* has an invalid address.

4703       **EPERM**

4704       The process does not have appropriate privileges.

4705       **EINVAL**

4706       *size* is greater than **NGROUPS\_MAX**.

**sethostname****Name**

4707       **sethostname** — set host name

**Synopsis**

```
4708 #include <unistd.h>
4709 #include <sys/param.h.h>
```

```
4710     #include <sys/utsname.h>
4711     int sethostname(const char * name, size_t len);
```

## Description

4712 If the process has appropriate privileges, the `sethostname()` function shall change  
 4713 the host name for the current machine. The `name` shall point to a null-terminated  
 4714 string of at most `len` bytes that holds the new hostname.

4715 If the symbol `HOST_NAME_MAX` is defined, or if `sysconf(_SC_HOST_NAME_MAX)()`  
 4716 returns a value greater than 0, this value shall represent the maximum length of the  
 4717 new hostname. Otherwise, if the symbol `MAXHOSTLEN` is defined, this value shall  
 4718 represent the maximum length for the new hostname. If none of these values are  
 4719 defined, the maximum length shall be the size of the `nodename` field of the `utsname`  
 4720 structure.

## Return Value

4721 On success, 0 is returned. On error, -1 is returned and the global variable `errno` is set  
 4722 appropriately.

## Errors

4723 **EINVAL**

4724     `len` is negative or larger than the maximum allowed size.

4725 **EPERM**

4726     the process did not have appropriate privilege.

4727 **EFAULT**

4728     `name` is an invalid address.

## Rationale

4729 ISO POSIX (2003) guarantees that:

4730     Maximum length of a host name (not including the terminating null) as returned from  
 4731     the `gethostname()` function shall be at least 255 bytes.

4732 The glibc C library does not currently define `HOST_NAME_MAX`, and although it  
 4733 provides the name `_SC_HOST_NAME_MAX` a call to `sysconf()` returns -1 and does not  
 4734 alter `errno` in this case (indicating that there is no restriction on the hostname  
 4735 length). However, the glibc manual indicates that some implementations may have  
 4736 `MAXHOSTNAMELEN` as a means of detecting the maximum length, while the Linux  
 4737 kernel at release 2.4 and 2.6 stores this hostname in the `utsname` structure. While the  
 4738 glibc manual suggests simply shortening the name until `sethostname()` succeeds,  
 4739 the LSB requires that one of the first four mechanisms works. Future versions of  
 4740 glibc may provide a more reasonable result from `sysconf(_SC_HOST_NAME_MAX)`.

## setsockopt

### Name

4741 setsockopt — set socket options

### Synopsis

4742 #include <sys/socket.h>

```
4743 #include <netinet/ip.h>
4744 int setsockopt(int socket, int level, int option_name, const void *
4745 option_value, socklen_t option_len);
```

## Description

4746 The `setsockopt()` function shall behave as specified in ISO POSIX (2003), with the  
 4747 following extensions.

### IP Protocol Level Options

4749 If the `level` parameter is  `IPPROTO_IP`, the following values shall be supported for  
 4750 `option_name` (see RFC 791:Internet Protocol for further details):

4751 `IP_OPTIONS`

4752 Set the Internet Protocol options sent with every packet from this socket. The  
 4753 `option_value` shall point to a memory buffer containing the options and  
 4754 `option_len` shall contain the size in bytes of that buffer. For IPv4, the  
 4755 maximum length of options is 40 bytes.

4756 `IP_TOS`

4757 Set the Type of Service flags to use when sending packets with this socket. The  
 4758 `option_value` shall point to a value containing the type of service value. The  
 4759 least significant two bits of the value shall contain the new Type of Service  
 4760 indicator. Use of other bits in the value is unspecified. The `option_len`  
 4761 parameter shall hold the size, in bytes, of the buffer referred to by  
 4762 `option_value`.

4763 `IP_TTL`

4764 Set the current unicast Internet Protocol Time To Live value used when sending  
 4765 packets with this socket. The `option_value` shall point to a value containing the  
 4766 time to live value, which shall be between 1 and 255. The `option_len`  
 4767 parameter shall hold the size, in bytes, of the buffer referred to by  
 4768 `option_value`.

4769 `IP_MULTICAST_TTL`

4770 Sets the Time To Live value of outgoing multicast packets for this socket.  
 4771 `optval` shall point to an integer which contains the new TTL value. If the new  
 4772 TTL value is -1, the implementation should use an unspecified default TTL  
 4773 value. If the new TTL value is out of the range of acceptable values (0-255),  
 4774 `setsockopt()` shall return -1 and set `errno` to indicate the error.

4775 `IP_MULTICAST_LOOP`

4776 Sets a boolean flag indicating whether multicast packets originating locally  
 4777 should be looped back to the local sockets. `optval` shall point to an integer  
 4778 which contains the new flag value.

4779 `IP_ADD_MEMBERSHIP`

4780 Join a multicast group. `optval` shall point to a `ip_mreq` structure. Before calling,  
 4781 the caller should fill in the `imr_multiaddr` field with the multicast group  
 4782 address and the `imr_interface` field with the address of the local interface. If  
 4783 `imr_address` is set to INADDR\_ANY, then an appropriate interface is chosen  
 4784 by the system.

4785           **IP\_DROP\_MEMBERSHIP**  
 4786           Leave a multicast group. *optval* shall point to a *ip\_mreq* structure containing  
 4787           the same values as were used with **IP\_ADD\_MEMBERSHIP**.  
 4788           **IP\_MULTICAST\_IF**  
 4789           Set the local device for a multicast socket. *optval* shall point to a *ip\_mreq*  
 4790           structure initialized in the same manner as with **IP\_ADD\_MEMBERSHIP**.  
 4791           The *ip\_mreq* structure contains two `struct in_addr` fields: *imr\_multiaddr* and  
 4792           *imr\_address*.

## Return Value

4793           On success, 0 is returned. On error, -1 is returned and the global variable `errno` is set  
 4794           appropriately.

## Errors

4795           As defined in ISO POSIX (2003).

## setutent

### Name

4796           **setutent** — access user accounting database entries

### Synopsis

```
4797 #include <utmp.h>
4798 void setutent(void);
```

### Description

4799           The **setutent()** function shall reset the user accounting database such that the next  
 4800           call to **getutent()** shall return the first record in the database. It is recommended to  
 4801           call it before any of the other functions that operate on the user accounting databases  
 4802           (e.g. **getutent()**)

### Return Value

4803           None.

## sigandset

### Name

4804        sigandset — build a new signal set by combining the two input sets using logical  
 4805        AND

### Synopsis

4806        #include <signal.h>  
 4807        int sigandset(sigset\_t \* *set*, const sigset\_t \* *left*, const sigset\_t \* *right*);

### Description

4808        The `sigandset()` shall combine the two signal sets referenced by *left* and *right*,  
 4809        using a logical AND operation, and shall place the result in the location referenced  
 4810        by *set*. The resulting signal set shall contain only signals that are in both the set  
 4811        referenced by *left* and the set referenced by *right*.

### Return Value

4812        On success, `sigandset()` shall return 0. Otherwise, `sigandset()` shall return -1 and  
 4813        set `errno` to indicate the error.

### Errors

4814        EINVAL  
 4815        One or more of *set*, *left*, or *right* was a null pointer.

### See Also

4816        `sigorset()`

## sigisemptyset

### Name

4817        `sigisemptyset` — check for empty signal set

### Synopsis

4818        #include <signal.h>  
 4819        int sigisemptyset(const sigset\_t \* *set*);

### Description

4820        The `sigisemptyset()` function shall check for empty signal set referenced by *set*.

### Return Value

4821        The `sigisemptyset()` function shall return a positive non-zero value if the signal  
 4822        set referenced by *set* is empty, or zero if this set is empty. On error,  
 4823        `sigisemptyset()` shall return -1 and set `errno` to indicate the error.

### Errors

4824        EINVAL  
 4825        *set* is a null pointer.

## sigorset

### Name

4826        sigorset — build a new signal set by combining the two input sets using logical  
 4827        OR

### Synopsis

4828        #include <signal.h>  
 4829        int sigorset(sigset\_t \* *set*, const sigset\_t \* *left*, const sigset\_t \* *right*);

### Description

4830        The *sigorset()* shall combine the two signal sets referenced by *left* and *right*,  
 4831        using a logical OR operation, and shall place the result in the location referenced by  
 4832        *set*. The resulting signal set shall contain only signals that are in either the set  
 4833        referenced by *left* or the set referenced by *right*.

### Return Value

4834        On success, *sigorset()* shall return 0. Otherwise, *sigorset()* shall return -1 and set  
 4835        *errno* to indicate the error.

### Errors

4836        EINVAL  
 4837        One or more of *set*, *left*, or *right* was a null pointer.

### See Also

4838        *sigandset()*

## sigreturn

### Name

4839        *sigreturn* — return from signal handler and cleanup stack frame

### Synopsis

4840        int sigreturn(struct sigcontext \* *scp*);

### Description

4841        The *sigreturn()* function is used by the system to cleanup after a signal handler  
 4842        has returned. This function is not in the source standard; it is only in the binary  
 4843        standard.

### Return Value

4844        *sigreturn()* never returns.

## sscanf

### Name

4845      `sscanf` — convert formatted input

### Description

4846      The `scanf()` family of functions shall behave as described in ISO POSIX (2003),  
 4847      except as noted below.

### Differences

4848      The `%s`, `%S` and `%[` conversion specifiers shall accept an option length modifier `a`,  
 4849      which shall cause a memory buffer to be allocated to hold the string converted. In  
 4850      such a case, the argument corresponding to the conversion specifier should be a  
 4851      reference to a pointer value that will receive a pointer to the allocated buffer. If there  
 4852      is insufficient memory to allocate a buffer, the function may set `errno` to `ENOMEM`  
 4853      and a conversion error results.

4854      **Note:** This directly conflicts with the ISO C (1999) usage of `%a` as a conversion specifier  
 4855      for hexadecimal float values. While this conversion specifier should be supported, a  
 4856      format specifier such as "`%aseconds`" will have a different meaning on an LSB  
 4857      conforming system.

## stime

### Name

4858      `stime` — set time

### Synopsis

```
4859 #define _SVID_SOURCE
4860 #include <time.h>
4861 int stime(const time_t * t);
```

### Description

4862      If the process has appropriate privilege, the `stime()` function shall set the system's  
 4863      idea of the time and date. Time, referenced by `t`, is measured in seconds from the  
 4864      epoch (defined in ISO POSIX (2003) as 00:00:00 UTC January 1, 1970).

### Return Value

4865      On success, `stime()` shall return 0. Otherwise, `stime()` shall return -1 and `errno`  
 4866      shall be set to indicate the error.

### Errors

4867      `EPERM`

4868      The process does not have appropriate privilege.

4869      `EINVAL`

4870      `t` is a null pointer.

## stpcpy

### Name

4871 stpcpy — copy a string returning a pointer to its end

### Synopsis

4872 #include <string.h>  
4873 char \* stpcpy(char \* restrict dest, const char \* restrict src);

### Description

4874 The `stpcpy()` function shall copy the string pointed to by `src` (including the  
4875 terminating null character) to the array pointed to by `dest`. The strings may not  
4876 overlap, and the destination string `dest` shall be large enough to receive the copy.

### Return Value

4877 `stpcpy()` returns a pointer to the end of the string `dest` (that is, the address of the  
4878 terminating null character) rather than the beginning.

### Example

4879 This program uses `stpcpy()` to concatenate `foo` and `bar` to produce `foobar`, which  
4880 it then prints.

```
4881 #include <string.h>
4882
4883     int
4884     main (void)
4885     {
4886         char buffer[256];
4887         char *to = buffer;
4888         to = stpcpy (to, "foo");
4889         to = stpcpy (to, "bar");
4890         printf ("%s\n", buffer);
4891     }
```

## stpncpy

### Name

4892        `stpncpy` — copy a fixed-size string, returning a pointer to its end

### Synopsis

4893        `#include <string.h>`  
 4894        `char * stpncpy(char * restrict dest, const char * restrict src, size_t n);`

### Description

4895        The `stpncpy()` function shall copy at most *n* characters from the string pointed to by  
 4896        *src*, including the terminating null character, to the array pointed to by *dest*.  
 4897        Exactly *n* characters are written at *dest*. If the length `strlen(src)` is smaller than  
 4898        *n*, the remaining characters in *dest* are filled with '\0' characters. If the length  
 4899        `strlen(src)` is greater than or equal to *n*, *dest* will not be null terminated.

4900        The strings may not overlap.

4901        The programmer shall ensure that there is room for at least *n* characters at *dest*.

### Return Value

4902        The `stpncpy()` function shall return a pointer to the terminating NULL in *dest*, or,  
 4903        if *dest* is not NULL-terminated, *dest* + *n*.

## strcasestr

### Name

4904        `strcasestr` — locate a substring ignoring case

### Synopsis

4905        `#include <string.h>`  
 4906        `char * strcasestr(const char * s1, const char * s2);`

### Description

4907        The `strcasestr()` shall behave as `strstr()`, except that it shall ignore the case of  
 4908        both strings. The `strcasestr()` function shall be locale aware; that is `strcasestr()`  
 4909        shall behave as if both strings had been converted to lower case in the current locale  
 4910        before the comparison is performed.

### Return Value

4911        Upon successful completion, `strcasestr()` shall return a pointer to the located  
 4912        string or a null pointer if the string is not found. If *s2* points to a string with zero  
 4913        length, the function shall return *s1*.

**strerror\_r****Name**

4914           **strerror\_r** — reentrant version of strerror

**Synopsis**

4915           **#include <string.h>**  
 4916           **char \* strerror\_r(int errnum, char \* buf, size\_t buflen);**

**Description**

4917           The **strerror\_r()** shall behave as specified in ISO POSIX (2003), except as  
 4918           described below.

**Returns String, not Error Value**

4920           The **strerror\_r()** function shall return a pointer to the string corresponding to  
 4921           *errno*. The returned pointer may point within the buffer *buf* (at most *buflen* bytes).

**Return Value**

4922           On success, **strerror\_r()** shall return a pointer to the generated message string  
 4923           (determined by the setting of the LC\_MESSAGES category in the current locale).  
 4924           Otherwise, **strerror\_r()** shall return the string corresponding to "Unknown error".

**strndup****Name**

4925           **strndup** — return a malloc'd copy of at most the specified number of bytes of a  
 4926           string

**Synopsis**

4927           **#include <string.h>**  
 4928           **char \* strndup(const char \* string, size\_t n);**

**Description**

4929           The **strndup()** function shall return a malloc()'d copy of at most *n* bytes of *string*.  
 4930           The resultant string shall be terminated even if no NULL terminator appears before  
 4931           *string+n*.

**Return Value**

4932           On success, **strndup()** shall return a pointer to a newly allocated block of memory  
 4933           containing a copy of at most *n* bytes of *string*. Otherwise, **strndup()** shall return  
 4934           NULL and set *errno* to indicate the error.

**Errors**

4935           **ENOMEM**

4936           Insufficient memory available.

## strnlen

### Name

4937 `strnlen` – determine the length of a fixed-size string

### Synopsis

4938 `#include <string.h>`  
 4939 `size_t strnlen(const char * s, size_t maxlen);`

### Description

4940 `strnlen()` returns the number of characters in the string *s*, not including the  
 4941 terminating \0 character, but at most *maxlen*. In doing this, `strnlen()` looks only at  
 4942 the first *maxlen* characters at *s* and never beyond *s + maxlen*.

### Return Value

4943 `strnlen()` returns `strlen(s)`, if that is less than *maxlen*, or *maxlen* if there is no \0  
 4944 character among the first *maxlen* characters pointed to by *s*.

## strftime

### Name

4945 `strftime` – parse a time string

### Description

4946 The `strftime()` shall behave as specified in the *ISO POSIX (2003)* with differences  
 4947 as listed below.

### Number of leading zeroes may be limited

The *ISO POSIX (2003)* specifies fields for which "leading zeros are permitted but not required"; however, applications shall not expect to be able to supply more leading zeroes for these fields than would be implied by the range of the field. Implementations may choose to either match an input with excess leading zeroes, or treat this as a non-matching input. For example, %j has a range of 001 to 366, so 0, 00, 000, 001, and 045 are acceptable inputs, but inputs such as 0000, 0366 and the like are not.

### Rationale

4956 *glibc* developers consider it appropriate behavior to forbid excess leading zeroes.  
 4957 When trying to parse a given input against several format strings, forbidding excess  
 4958 leading zeroes could be helpful. For example, if one matches 0011-12-26  
 4959 against %m-%d-%Y and then against %Y-%m-%d, it seems useful for the first match to  
 4960 fail, as it would be perverse to parse that date as November 12, year 26. The second  
 4961 pattern parses it as December 26, year 11.

4962 The *ISO POSIX (2003)* is not explicit that an unlimited number of leading zeroes are  
 4963 required, although it may imply this. The LSB explicitly allows implementations to  
 4964 have either behavior. Future versions of this standard may require implementations  
 4965 to forbid excess leading zeroes.

4966 An Interpretation Request is currently pending against ISO POSIX (2003) for this  
 4967 matter.

## strsep

### Name

4968 strsep — extract token from string

### Synopsis

```
4969 #include <string.h>
4970 char * strsep(char * * stringp, const char * delim);
```

### Description

4971 The `strsep()` function shall find the first token in the string referenced by the  
4972 pointer `stringp`, using the characters in `delim` as delimiters.

4973 If `stringp` is `NULL`, `strsep()` shall return `NULL` and do nothing else.

4974 If `stringp` is non-`NULL`, `strsep()` shall find the first token in the string referenced  
4975 by `stringp`, where tokens are delimited by characters in the string `delim`. This token  
4976 shall be terminated with a `\0` character by overwriting the delimiter, and `stringp`  
4977 shall be updated to point past the token. In case no delimiter was found, the token is  
4978 taken to be the entire string referenced by `stringp`, and the location referenced by  
4979 `stringp` is made `NULL`.

### Return Value

4980 `strsep()` shall return a pointer to the beginning of the token.

### Notes

4981 The `strsep()` function was introduced as a replacement for `strtok()`, since the  
4982 latter cannot handle empty fields. However, `strtok()` conforms to ISO C (1999) and  
4983 to ISO POSIX (2003) and hence is more portable.

### See Also

4984 `strtok()`, `strtok_r()`.

## strsignal

### Name

4985 `strsignal` — return string describing signal

### Synopsis

```
4986 #define __GNU_SOURCE
```

```
4987     #include <string.h>
4988     char * strsignal(int sig);
4989
4990     extern const char * const sys_siglist[];
```

## Description

4990 The `strsignal()` function shall return a pointer to a string describing the signal  
4991 number *sig*. The string can only be used until the next call to `strsignal()`.  
4992 The array `sys_siglist` holds the signal description strings indexed by signal  
4993 number. This array should not be accessed directly by applications.

## Return Value

4994 If *sig* is a valid signal number, `strsignal()` shall return a pointer to the  
4995 appropriate description string. Otherwise, `strsignal()` shall return either a pointer  
4996 to the string "unknown signal", or a null pointer.  
4997 Although the function is not declared as returning a pointer to a constant character  
4998 string, applications shall not modify the returned string.

## strtouq

### Name

4999 `strtouq` — convert string value to a long or quad\_t integer

### Synopsis

```
5000 #include <sys/types.h>
5001 #include <stdlib.h>
```

```
5002 #include <limits.h>
5003 long long strtouq(const char * nptr, char ** endptr, int base);
```

## Description

5004  
 5005     `strtouq()` converts the string *nptr* to a quadt value. The conversion is done  
 5006     according to the given base, which shall be between 2 and 36 inclusive, or be the  
       special value 0.  
 5007  
 5008     *nptr* may begin with an arbitrary amount of white space (as determined by  
 5009     *isspace()*), followed by a single optional + or - sign character. If *base* is 0 or 16, the  
 5010     string may then include a 0x prefix, and the number will be read in base 16;  
 5011     otherwise, a 0 base is taken as 10 (decimal), unless the next character is 0, in which  
       case it is taken as 8 (octal).

5012  
 5013     The remainder of the string is converted to a long value in the obvious manner,  
 5014     stopping at the first character which is not a valid digit in the given base. (In bases  
 5015     above 10, the letter A in either upper or lower case represents 10, B represents 11, and  
       so forth, with Z representing 35.)

## Return Value

5016     `strtouq()` returns the result of the conversion, unless the value would underflow or  
 5017     overflow. If an underflow occurs, `strtouq()` returns QUAD\_MIN. If an overflow occurs,  
 5018     `strtouq()` returns QUAD\_MAX. In both cases, the global variable `errno` is set to  
 5019     ERANGE.

## Errors

5020     ERANGE  
 5021         The given string was out of range; the value converted has been clamped.

## **strtouq**

### Name

5022     `strtouq` — convert a string to an unsigned long long

### Synopsis

```
5023 #include <sys/types.h>
5024 #include <stdlib.h>
```

```
5025     #include <limits.h>
5026     unsigned long long strtouq(const char * nptr, char * * endptr, int base);
```

## Description

5027       `strtouq()` converts the string *nptr* to an unsigned long long value. The conversion  
 5028       is done according to the given base, which shall be between 2 and 36 inclusive, or be  
 5029       the special value 0.

5030       *nptr* may begin with an arbitrary amount of white space (as determined by  
 5031       *isspace()*), followed by a single optional + or - sign character. If *base* is 0 or 16, the  
 5032       string may then include a 0x prefix, and the number will be read in base 16;  
 5033       otherwise, a 0 base is taken as 10 (decimal), unless the next character is 0, in which  
 5034       case it is taken as 8 (octal).

5035       The remainder of the string is converted to an unsigned long value in the obvious  
 5036       manner, stopping at the end of the string or at the first character that does not  
 5037       produce a valid digit in the given base. (In bases above 10, the letter A in either upper  
 5038       or lower case represents 10, B represents 11, and so forth, with Z representing 35.)

## Return Value

5039       On success, `strtouq()` returns either the result of the conversion or, if there was a  
 5040       leading minus sign, the negation of the result of the conversion, unless the original  
 5041       (non-negated) value would overflow. In the case of an overflow the function returns  
 5042       `UQUAD_MAX` and the global variable `errno` is set to ERANGE.

## Errors

5043       ERANGE

5044       The given string was out of range; the value converted has been clamped.

## **svc\_register**

### Name

5045       **svc\_register** — register Remote Procedure Call interface

### Synopsis

```
5046     #include <rpc/rpc.h>
5047     bool_t svc_register(SVCXPRT * xprt, rpcprog_t progrnum, rpcvers_t versnum,
5048     __dispatch_fn_t dispatch, rpcprot_t protocol);
```

## Description

5049       The `svc_register()` function shall associate the program identified by *progrnum* at  
 5050       version *versnum* with the service dispatch procedure, *dispatch*. If *protocol* is zero,  
 5051       the service is not registered with the portmap service. If *protocol* is non-zero, then a  
 5052       mapping of the triple [*progrnum*, *versnum*, *protocol*] to *xprt->xp\_port* is  
 5053       established with the local portmap service. The procedure *dispatch* has the  
 5054       following form:

```
5055     int dispatch(struct svc_req * request, SVCXPRT * xprt);
```

## Return Value

5056       **svc\_register()** returns 1 if it succeeds, and zero otherwise.

**svc\_run****Name**

5057           svc\_run — waits for RPC requests to arrive and calls service procedure

**Synopsis**

5058           #include <rpc/svc.h>  
5059           void svc\_run(void);

**Description**

5060           The `svc_run()` function shall wait for RPC requests to arrive, read and unpack each  
5061           request, and dispatch it to the appropriate registered handler. Under normal  
5062           conditions, `svc_run()` shall not return; it shall only return if serious errors occur  
5063           that prevent further processing.

**svc\_sendreply****Name**

5064           svc\_sendreply — called by RPC service's dispatch routine

**Synopsis**

5065           bool\_t svc\_sendreply(SVCXPRT \*xprt, xdrproc\_t outproc, caddr\_t out);

**Description**

5066           Called by an RPC service's dispatch routine to send the results of a remote  
5067           procedure call. The parameter `xprt` is the request's associated transport handle;  
5068           `outproc` is the XDR routine which is used to encode the results; and `out` is the  
5069           address of the results. This routine returns one if it succeeds, zero other-wise.

**svctcp\_create****Name**

5070 svctcp\_create — create a TCP/IP-based RPC service transport

**Synopsis**

```
5071 #include <rpc/rpc.h>
5072 SVCXPRT * svctcp_create(int sock, u_int send_buf_size, u_int recv_buf_size);
```

**Description**

5073 svctcp\_create( ) creates a TCP/IP-based RPC service transport, to which it returns  
 5074 a pointer. The transport is associated with the socket *sock*, which may be  
 5075 `RPC_ANYSOCK`, in which case a new socket is created. If the socket is not bound to a  
 5076 local TCP port, then this routine binds it to an arbitrary port. Upon completion,  
 5077 `xprt->xp_sock` is the transport's socket descriptor, and `xprt->xp_port` is the  
 5078 transport's port number. Since TCP-based RPC uses buffered I/O, users may specify  
 5079 the size of buffers; values of zero choose suitable defaults.

**Return Value**

5080 `svctcp_create( )` returns `NULL` if it fails, or a pointer to the RPC service transport  
 5081 otherwise.

**svcudp\_create****Name**

5082 `svcudp_create — create a UDP-based RPC service transport`

**Synopsis**

```
5083 SVCXPRT *
5084 svcudp_create(int sock);
```

**Description**

5085 This call is equivalent to `svcudp_bufcreate( sock, sz, sz)` for some default size  
 5086 `sz`.

## swscanf

### Name

5087      `swscanf` — convert formatted input

### Description

5088      The `scanf()` family of functions shall behave as described in ISO POSIX (2003),  
5089      except as noted below.

### Differences

5090      The `%s`, `%S` and `%[` conversion specifiers shall accept an option length modifier `a`,  
5091      which shall cause a memory buffer to be allocated to hold the string converted. In  
5092      such a case, the argument corresponding to the conversion specifier should be a  
5093      reference to a pointer value that will receive a pointer to the allocated buffer. If there  
5094      is insufficient memory to allocate a buffer, the function may set `errno` to `ENOMEM`  
5095      and a conversion error results.

5096      **Note:** This directly conflicts with the ISO C (1999) usage of `%a` as a conversion specifier  
5097      for hexadecimal float values. While this conversion specifier should be supported, a  
5098      format specifier such as "`%aseconds`" will have a different meaning on an LSB  
5099      conforming system.

## system

### Name

5100 system — execute a shell command

### Synopsis

5101 #include <stdlib.h>  
5102 int system(const char \* string);

### Description

5103 The `system()` function shall behave as described in ISO POSIX (2003).

### Notes

5104 The fact that `system()` ignores interrupts is often not what a program wants. ISO  
5105 POSIX (2003) describes some of the consequences; an additional consequence is that  
5106 a program calling `system()` from a loop cannot be reliably interrupted. Many  
5107 programs will want to use the `exec()` family of functions instead.

5108 Do not use `system()` from a program with `suid` or `sgid` privileges, because  
5109 unexpected values for some environment variables might be used to subvert system  
5110 integrity. Use the `exec()` family of functions instead, but not `execlp()` or `execvp()`.  
5111 `system()` will not, in fact, work properly from programs with `suid` or `sgid`  
5112 privileges on systems on which `/bin/sh` is **bash** version 2, since **bash** 2 drops  
5113 privileges on startup. (Debian uses a modified **bash** which does not do this when  
5114 invoked as **sh**.)

5115 The check for the availability of `/bin/sh` is not actually performed; it is always  
5116 assumed to be available. ISO C (1999) specifies the check, but ISO POSIX (2003)  
5117 specifies that the return shall always be nonzero, since a system without the shell is  
5118 not conforming, and it is this that is implemented.

5119 It is possible for the shell command to return 127, so that code is not a sure indication  
5120 that the `execve()` call failed; check the global variable `errno` to make sure.

## textdomain

### Name

5121           textdomain — set the current default message domain

### Synopsis

```
5122 #include <libintl.h>
5123 char * textdomain(const char * domainname);
```

### Description

5124           The `textdomain()` function shall set the current default message domain to  
 5125           *domainname*. Subsequent calls to `gettext()` and `ngettext()` use the default  
 5126           message domain.

5127           If *domainname* is NULL, the default message domain shall not be altered.

5128           If *domainname* is "", `textdomain()` shall reset the default domain to the system  
 5129           default of "messages".

### Return

5130           On success, `textdomain()` shall return the currently selected domain. Otherwise, a  
 5131           null pointer shall be returned, and `errno` is set to indicate the error.

### Errors

5132           ENOMEM

5133           Insufficient memory available.

## unlink

### Name

5134           unlink — remove a directory entry

### Synopsis

```
5135 int unlink(const char * path);
```

### Description

5136           `unlink()` is as specified in ISO POSIX (2003), but with differences as listed below.

5137           See also Section 18.1, Additional behaviors: `unlink/link` on directory.

5138           **May return EISDIR on directories**

5139           If *path* specifies a directory, the implementation may return EISDIR instead of  
 5140           EPERM as specified by ISO POSIX (2003).

5141           **Rationale:** The Linux kernel has deliberately chosen EISDIR for this case and does not  
 5142           expect to change.

## uselocale

### Name

5143        uselocale — set locale for thread

### Synopsis

5144        #include <locale.h>  
5145        locale\_t uselocale(locale\_t newloc);

### Description

5146        The uselocale() function shall set the locale for the calling thread to the locale  
5147        specified by *newloc*.

5148        If *newloc* is the value LC\_GLOBAL\_LOCALE, the thread's locale shall be set to the  
5149        process current global locale, as set by setlocale(). If *newloc* is NULL, the thread's  
5150        locale is not altered.

### Return Value

5151        The uselocale() function shall return the previous locale, or LC\_GLOBAL\_LOCALE if  
5152        the thread local locale has not been previously set.

### Errors

5153        None defined.

### See Also

5154        setlocale(), freelocale(), duplocale(), newlocale()

## utmpname

### Name

5155 utmpname — set user accounting database

### Synopsis

```
5156 #include <utmp.h>
5157 int utmpname(const char * dbname);
```

### Description

5158 The `utmpname()` function shall cause the user accounting database used by the  
 5159 `getutent()`, `getutent_r()`, `getutxent()`, `getutxid()`, `getutxline()`, and  
 5160 `pututxline()` functions to be that named by `dbname`, instead of the system default  
 5161 database. See Section 16.3 for further information.

5162 **Note:** The LSB does not specify the format of the user accounting database, nor the  
 5163 names of the file or files that may contain it.

### Return Value

5164 None.

### Errors

5165 None defined.

## vasprintf

### Name

5166 `vasprintf` — write formatted output to a dynamically allocated string

### Synopsis

```
5167 #include <stdarg.h>
5168 #include <stdio.h>
5169 int vasprintf(char * * restrict ptr, const char * restrict format, va_list
5170 arg);
```

### Description

5171 The `vasprintf()` function shall write formatted output to a dynamically allocated  
 5172 string, and store the address of that string in the location referenced by `ptr`. It shall  
 5173 behave as `asprintf()`, except that instead of being called with a variable number of  
 5174 arguments, it is called with an argument list as defined by `<stdarg.h>`.

### Return Value

5175 Refer to `fprintf()`.

### Errors

5176 Refer to `fprintf()`.

## vdprintf

### Name

5177 vdprintf — write formatted output to a file descriptor

### Synopsis

```
5178 #include <stdio.h>
5179 int vdprintf(int fd, const char * restrict format, va_list arg);
```

### Description

5180 The `vdprintf()` function shall behave as `vfprintf()`, except that `vdprintf()` shall  
 5181 write output to the file associated with the file descriptor specified by the `fd`  
 5182 argument, rather than place output on a stream (as defined by ISO POSIX (2003)).

### Return Value

5183 Refer to `fprintf()`.

### Errors

5184 Refer to `fprintf()`.

## verrx

### Name

5185 `verrx` — display formatted error message and exit

### Synopsis

```
5186 #include <stdarg.h>
5187 #include <err.h>
5188 void verrx(int eval, const char * fmt, va_list args);
```

### Description

5189 The `verrx()` shall behave as `errx()` except that instead of being called with a  
 5190 variable number of arguments, it is called with an argument list as defined by  
 5191 `<stdarg.h>`.

5192 `verrx()` does not return, but exits with the value of `eval`.

### Return Value

5193 None.

### Errors

5194 None.

## vfscanf

### Name

5195 vfscanf — convert formatted input

### Description

5196 The `scanf()` family of functions shall behave as described in ISO POSIX (2003),  
 5197 except as noted below.

### Differences

5198 The `%s`, `%S` and `%[` conversion specifiers shall accept an option length modifier `a`,  
 5199 which shall cause a memory buffer to be allocated to hold the string converted. In  
 5200 such a case, the argument corresponding to the conversion specifier should be a  
 5201 reference to a pointer value that will receive a pointer to the allocated buffer. If there  
 5202 is insufficient memory to allocate a buffer, the function may set `errno` to `ENOMEM`  
 5203 and a conversion error results.

5204 **Note:** This directly conflicts with the ISO C (1999) usage of `%a` as a conversion specifier  
 5205 for hexadecimal float values. While this conversion specifier should be supported, a  
 5206 format specifier such as "`%aseconds`" will have a different meaning on an LSB  
 5207 conforming system.

## vfwscanf

### Name

5208 vfwscanf — convert formatted input

### Description

5209 The `scanf()` family of functions shall behave as described in ISO POSIX (2003),  
 5210 except as noted below.

### Differences

5211 The `%s`, `%S` and `%[` conversion specifiers shall accept an option length modifier `a`,  
 5212 which shall cause a memory buffer to be allocated to hold the string converted. In  
 5213 such a case, the argument corresponding to the conversion specifier should be a  
 5214 reference to a pointer value that will receive a pointer to the allocated buffer. If there  
 5215 is insufficient memory to allocate a buffer, the function may set `errno` to `ENOMEM`  
 5216 and a conversion error results.

5217 **Note:** This directly conflicts with the ISO C (1999) usage of `%a` as a conversion specifier  
 5218 for hexadecimal float values. While this conversion specifier should be supported, a  
 5219 format specifier such as "`%aseconds`" will have a different meaning on an LSB  
 5220 conforming system.

**vscanf****Name**

5221      vscanf — convert formatted input

**Description**

5222      The `scanf()` family of functions shall behave as described in ISO POSIX (2003),  
 5223      except as noted below.

**Differences**

5224      The `%s`, `%S` and `%[` conversion specifiers shall accept an option length modifier `a`,  
 5225      which shall cause a memory buffer to be allocated to hold the string converted. In  
 5226      such a case, the argument corresponding to the conversion specifier should be a  
 5227      reference to a pointer value that will receive a pointer to the allocated buffer. If there  
 5228      is insufficient memory to allocate a buffer, the function may set `errno` to `ENOMEM`  
 5229      and a conversion error results.

5230      **Note:** This directly conflicts with the ISO C (1999) usage of `%a` as a conversion specifier  
 5231      for hexadecimal float values. While this conversion specifier should be supported, a  
 5232      format specifier such as "`%aseconds`" will have a different meaning on an LSB  
 5233      conforming system.

**vsscanf****Name**

5234      vsscanf — convert formatted input

**Description**

5235      The `scanf()` family of functions shall behave as described in ISO POSIX (2003),  
 5236      except as noted below.

**Differences**

5237      The `%s`, `%S` and `%[` conversion specifiers shall accept an option length modifier `a`,  
 5238      which shall cause a memory buffer to be allocated to hold the string converted. In  
 5239      such a case, the argument corresponding to the conversion specifier should be a  
 5240      reference to a pointer value that will receive a pointer to the allocated buffer. If there  
 5241      is insufficient memory to allocate a buffer, the function may set `errno` to `ENOMEM`  
 5242      and a conversion error results.

5243      **Note:** This directly conflicts with the ISO C (1999) usage of `%a` as a conversion specifier  
 5244      for hexadecimal float values. While this conversion specifier should be supported, a  
 5245      format specifier such as "`%aseconds`" will have a different meaning on an LSB  
 5246      conforming system.

## vswscanf

### Name

5247 vswscanf — convert formatted input

### Description

5248 The `scanf()` family of functions shall behave as described in ISO POSIX (2003),  
 5249 except as noted below.

### Differences

5250 The %s, %S and %[ conversion specifiers shall accept an option length modifier a,  
 5251 which shall cause a memory buffer to be allocated to hold the string converted. In  
 5252 such a case, the argument corresponding to the conversion specifier should be a  
 5253 reference to a pointer value that will receive a pointer to the allocated buffer. If there  
 5254 is insufficient memory to allocate a buffer, the function may set `errno` to ENOMEM  
 5255 and a conversion error results.

5256 **Note:** This directly conflicts with the ISO C (1999) usage of %a as a conversion specifier  
 5257 for hexadecimal float values. While this conversion specifier should be supported, a  
 5258 format specifier such as "%aseconds" will have a different meaning on an LSB  
 5259 conforming system.

## vsyslog

### Name

5260 vsyslog — log to system log

### Synopsis

```
5261 #include <stdarg.h>
5262 #include <syslog.h>
5263 void vsyslog(int priority, char * message, va_list arglist);
```

### Description

5264 The `vsyslog()` function is identical to `syslog()` as specified in ISO POSIX (2003),  
 5265 except that `arglist` (as defined by `stdarg.h`) replaces the variable number of  
 5266 arguments.

**vwscanf****Name**

5267        `vwscanf` — convert formatted input

**Description**

5268        The `scanf()` family of functions shall behave as described in ISO POSIX (2003),  
 5269        except as noted below.

**Differences**

5270        The `%s`, `%S` and `%[` conversion specifiers shall accept an option length modifier `a`,  
 5271        which shall cause a memory buffer to be allocated to hold the string converted. In  
 5272        such a case, the argument corresponding to the conversion specifier should be a  
 5273        reference to a pointer value that will receive a pointer to the allocated buffer. If there  
 5274        is insufficient memory to allocate a buffer, the function may set `errno` to `ENOMEM`  
 5275        and a conversion error results.

5276        **Note:** This directly conflicts with the ISO C (1999) usage of `%a` as a conversion specifier  
 5277        for hexadecimal float values. While this conversion specifier should be supported, a  
 5278        format specifier such as "`%aseconds`" will have a different meaning on an LSB  
 5279        conforming system.

**wait4****Name**

5280        `wait4` — wait for process termination, BSD style

**Synopsis**

```
5281 #include <sys/types.h>
5282 #include <sys/resource.h>
```

```
5283 #include <sys/wait.h>
5284 pid_t wait4(pid_t pid, int * status, int options, struct rusage * rusage);
```

## Description

5285 `wait4()` suspends execution of the current process until a child (as specified by *pid*)  
 5286 has exited, or until a signal is delivered whose action is to terminate the current  
 5287 process or to call a signal handling function. If a child (as requested by *pid*) has  
 5288 already exited by the time of the call (a so-called "zombie" process), the function  
 5289 returns immediately. Any system resources used by the child are freed.

5290 The value of *pid* can be one of:

5291 < -1

5292     wait for any child process whose process group ID is equal to the absolute value  
 5293 of *pid*.

5294 -1

5295     wait for any child process; this is equivalent to calling `wait3()`.

5296 0

5297     wait for any child process whose process group ID is equal to that of the calling  
 5298 process.

5299 > 0

5300     wait for the child whose process ID is equal to the value of *pid*.

5301 The value of *options* is a bitwise or of zero or more of the following constants:

5302 WNOHANG

5303     return immediately if no child is there to be waited for.

5304 WUNTRACED

5305     return for children that are stopped, and whose status has not been reported.

5306 If *status* is not NULL, `wait4()` stores status information in the location *status*. This  
 5307 status can be evaluated with the following macros:

5308 **Note:** These macros take the *status* value (an *int*) as an argument -- not a pointer to the  
 5309 value!

5310 WIFEXITED(status)

5311     is nonzero if the child exited normally.

5312 WEXITSTATUS(status)

5313     evaluates to the least significant eight bits of the return code of the child that  
 5314 terminated, which may have been set as the argument to a call to `exit()` or as  
 5315 the argument for a return statement in the main program. This macro can only  
 5316 be evaluated if `WIFEXITED()` returned nonzero.

5317 WIFSIGNALED(status)

5318     returns true if the child process exited because of a signal that was not caught.

5319 WTERMSIG(status)

5320                 returns the number of the signal that caused the child process to terminate. This  
 5321                 macro can only be evaluated if `WIFSIGNALED()` returned nonzero.

5322                 **WIFSTOPPED(status)**  
 5323                 returns true if the child process that caused the return is currently stopped; this  
 5324                 is only possible if the call was done using `WUNTRACED()`.

5325                 **WSTOPSIG(status)**  
 5326                 returns the number of the signal that caused the child to stop. This macro can  
 5327                 only be evaluated if `WIFSTOPPED()` returned nonzero.  
 5328                 If *rusage* is not NULL, the struct *rusage* (as defined in `sys/resource.h`) that it  
 5329                 points to will be filled with accounting information. See `getrusage()` for details.

## Return Value

5330                 On success, the process ID of the child that exited is returned. On error, -1 is  
 5331                 returned (in particular, when no unwaited-for child processes of the specified kind  
 5332                 exist), or 0 if `WNOHANG()` was used and no child was available yet. In the latter two  
 5333                 cases, the global variable `errno` is set appropriately.

## Errors

5334                 **ECHILD**  
 5335                 No unwaited-for child process as specified does exist.  
 5336                 **ERESTARTSYS**  
 5337                 A `WNOHANG()` was not set and an unblocked signal or a `SIGCHILD` was caught.  
 5338                 This error is returned by the system call. The library interface is not allowed to  
 5339                 return ERESTARTSYS, but will return EINTR.

## waitpid

### Name

5340                 **waitpid** — wait for child process

### Description

5341                 **waitpid()** is as specified in ISO POSIX (2003), but with differences as listed below.

### Need not support `WCONTINUED` or `WIFCONTINUED`

5343                 Implementations need not support the XSI optional functionality of `WCONTINUED()`  
 5344                 or `WIFCONTINUED()`.

**warn****Name**

5345 warn — formatted error messages

**Synopsis**

```
5346 #include <err.h>
5347 void warn(const char * fmt, ...);
```

**Description**

5348 The `warn()` function shall display a formatted error message on the standard error  
 5349 stream. The output shall consist of the last component of the program name, a colon  
 5350 character, and a space character. If `fmt` is non-NULL, it shall be used as a format  
 5351 string for the `printf()` family of functions, and the formatted message, a colon  
 5352 character, and a space are written to `stderr`. Finally, the error message string  
 5353 affiliated with the current value of the global variable `errno` shall be written to  
 5354 `stderr`, followed by a newline character.

**Return Value**

5355 None.

**Errors**

5356 None.

**warnx****Name**

5357 `warnx` — formatted error messages

**Synopsis**

```
5358 #include <err.h>
5359 void warnx(const char * fmt, ...);
```

**Description**

5360 The `warnx()` function shall display a formatted error message on the standard error  
 5361 stream. The last component of the program name, a colon character, and a space  
 5362 shall be output. If `fmt` is non-NULL, it shall be used as the format string for the  
 5363 `printf()` family of functions, and the formatted error message, a colon character,  
 5364 and a space shall be output. The output shall be followed by a newline character.

**Return Value**

5365 None.

**Errors**

5366 None.

## wcpncpy

### Name

5367 `wcpncpy` — copy a wide character string, returning a pointer to its end

### Synopsis

```
5368 #include <wchar.h>
5369 wchar_t * wcpncpy(wchar_t * dest, const wchar_t * src);
```

### Description

5370 `wcpncpy()` is the wide-character equivalent of `stpncpy()`. It copies the wide character  
5371 string `src`, including the terminating null wide character code, to the array `dest`.

5372 The strings may not overlap.

5373 The programmer shall ensure that there is room for at least `wcslen()(src)+1` wide  
5374 characters at `dest`.

### Return Value

5375 `wcpncpy()` returns a pointer to the end of the wide-character string `dest`, that is, a  
5376 pointer to the terminating null wide character code.

## wcpncpy

### Name

5377 `wcpncpy` — copy a fixed-size string of wide characters, returning a pointer to its end

### Synopsis

```
5378 #include <wchar.h>
5379 wchar_t * wcpncpy(wchar_t * dest, const wchar_t * src, size_t n);
```

### Description

5380 `wcpncpy()` is the wide-character equivalent of `stpncpy()`. It copies at most `n` wide  
5381 characters from the wide-character string `src`, including the terminating null wide  
5382 character code, to the array `dest`. Exactly `n` wide characters are written at `dest`. If the  
5383 length `wcslen()(src)` is smaller than `n`, the remaining wide characters in the array  
5384 `dest` are filled with null wide character codes. If the length `wcslen()(src)` is  
5385 greater than or equal to `n`, the string `dest` will not be terminated with a null wide  
5386 character code.

5387 The strings may not overlap.

5388 The programmer shall ensure that there is room for at least `n` wide characters at  
5389 `dest`.

### Return Value

5390 `wcpncpy()` returns a pointer to the wide character one past the last non-null wide  
5391 character written.

## wcscasecmp

### Name

5392        `wcscasecmp` — compare two wide-character strings, ignoring case

### Synopsis

```
5393 #include <wchar.h>
5394 int wcscasecmp(const wchar_t * s1, const wchar_t * s2);
```

### Description

5395        `wcscasecmp()` is the wide-character equivalent of `strcasecmp()`. It compares the  
 5396        wide-character string *s1* and the wide-character string *s2*, ignoring case differences  
 5397        (`toupper`, `towlower`).

### Return Value

5398        The `wcscasecmp()` function shall return 0 if the wide-character strings *s1* and *s2* are  
 5399        equal except for case distinctions. It shall return a positive integer if *s1* is greater  
 5400        than *s2*, ignoring case. It shall return a negative integer if *s1* is less than *s2*, ignoring  
 5401        case.

### Notes

5402        The behavior of `wcscasecmp()` depends upon the `LC_CTYPE` category of the current  
 5403        locale.

## wcsdup

### Name

5404        `wcsdup` — duplicate a wide-character string

### Synopsis

```
5405 #include <wchar.h>
5406 wchar_t * wcsdup(const wchar_t * s);
```

### Description

5407        `wcsdup()` is the wide-character equivalent of `strdup()`. It allocates and returns a  
 5408        new wide-character string whose initial contents is a duplicate of the wide-character  
 5409        string *s*.

5410        Memory for the new wide-character string is obtained with `malloc()`, and can be  
 5411        freed with `free()`.

### Return Value

5412        `wcsdup()` returns a pointer to the new wide-character string, or NULL if sufficient  
 5413        memory was not available.

## wcsncasecmp

### Name

5414        `wcsncasecmp` — compare two fixed-size wide-character strings, ignoring case

### Synopsis

5415        

```
#include <wchar.h>
5416      int wcsncasecmp(const wchar_t * s1, const wchar_t * s2, size_t n);
```

### Description

5417        `wcsncasecmp()` is the wide-character equivalent of `strncasecmp()`. It compares the  
 5418        wide-character string *s1* and the wide-character string *s2*, but at most *n* wide  
 5419        characters from each string, ignoring case differences (`toupper`, `towlower`).

### Return Value

5420        `wcsncasecmp()` returns 0 if the wide-character strings *s1* and *s2*, truncated to at most  
 5421        length *n*, are equal except for case distinctions. It returns a positive integer if  
 5422        truncated *s1* is greater than truncated *s2*, ignoring case. It returns a negative integer  
 5423        if truncated *s1* is smaller than truncated *s2*, ignoring case.

### Notes

5424        The behavior of `wcsncasecmp()` depends upon the `LC_CTYPE` category of the current  
 5425        locale.

## wcsnlen

### Name

5426        `wcsnlen` — determine the length of a fixed-size wide-character string

### Synopsis

5427        

```
#include <wchar.h>
5428      size_t wcsnlen(const wchar_t * s, size_t maxlen);
```

### Description

5429        `wcsnlen()` is the wide-character equivalent of `strnlen()`. It returns the number of  
 5430        wide-characters in the string *s*, not including the terminating null wide character  
 5431        code, but at most *maxlen*. In doing this, `wcsnlen()` looks only at the first *maxlen*  
 5432        wide-characters at *s* and never beyond *s + maxlen*.

### Return Value

5433        `wcsnlen()` returns `wcslen()(s)` if that is less than *maxlen*, or *maxlen* if there is no  
 5434        null wide character code among the first *maxlen* wide characters pointed to by *s*.

## wcsnrtombs

### Name

5435        `wcsnrtombs` — convert a wide character string to a multi-byte string

### Synopsis

```
5436 #include <wchar.h>
5437 size_t wcsnrtombs(char * dest, const wchar_t ** src, size_t nwc, size_t len,
5438 mbstate_t * ps);
```

### Description

5439        `wcsnrtombs`( ) is like `wcsrtombs`( ), except that the number of wide characters to be  
 5440 converted, starting at `src`, is limited to `nwc`.

5441        If `dest` is not a NULL pointer, `wcsnrtombs`( ) converts at most `nwc` wide characters  
 5442 from the wide-character string `src` to a multibyte string starting at `dest`. At most  
 5443 `len` bytes are written to `dest`. The state `ps` is updated.

5444        The conversion is effectively performed by repeatedly calling:

5445        `wcrtomb(dest, *src, ps)`

5446        as long as this call succeeds, and then incrementing `dest` by the number of bytes  
 5447 written and `src` by 1.

5448        The conversion can stop for three reasons:

- 5449        • A wide character has been encountered that cannot be represented as a multibyte  
   5450 sequence (according to the current locale). In this case `src` is left pointing to the  
   5451 invalid wide character, (`size_t`)(-1) is returned, and `errno` is set to `EILSEQ`.
- 5452        • `nws` wide characters have been converted without encountering a null wide  
   5453 character code, or the length limit forces a stop. In this case, `src` is left pointing to  
   5454 the next wide character to be converted, and the number bytes written to `dest` is  
   5455 returned.
- 5456        • The wide-character string has been completely converted, including the  
   5457 terminating null wide character code (which has the side effect of bringing back  
   5458 `ps` to the initial state). In this case, `src` is set to NULL, and the number of bytes  
   5459 written to `dest`, excluding the terminating null wide character code, is returned.

5460        If `dest` is NULL, `len` is ignored, and the conversion proceeds as above, except that the  
 5461 converted bytes are not written out to memory, and that no destination length limit  
 5462 exists.

5463        In both of the above cases, if `ps` is a NULL pointer, a static anonymous state only  
 5464 known to `wcsnrtombs`( ) is used instead.

5465        The programmer shall ensure that there is room for at least `len` bytes at `dest`.

### Return Value

5466        `wcsnrtombs`( ) returns the number of bytes that make up the converted part of  
 5467 multibyte sequence, not including the terminating null wide character code. If a  
 5468 wide character was encountered which could not be converted, (`size_t`)(-1) is  
 5469 returned, and the global variable `errno` set to `EILSEQ`.

### Notes

5470           The behavior of `wcsnrtombs()` depends on the `LC_CTYPE` category of the current  
 5471           locale.  
 5472           Passing `NULL` as `ps` is not multi-thread safe.

## **wcstoq**

### **Name**

5473           `wcstoq` — convert wide string to long long int representation

### **Synopsis**

```
5474 #include <wchar.h>
5475 long long int wcstoq(const wchar_t * restrict nptr, wchar_t ** restrict
5476 endptr, int base);
```

### **Description**

5477           The `wcstoq()` function shall convert the initial portion of the wide string `nptr` to  
 5478           long long int representation. It is identical to `wctoll()`.

### **Return Value**

5479           Refer to `wctoll()`.

### **Errors**

5480           Refer to `wctoll()`.

## **wcstouq**

### **Name**

5481           `wcstouq` — convert wide string to unsigned long long int representation

### **Synopsis**

```
5482 #include <wchar.h>
5483 unsigned long long wcstouq(const wchar_t * restrict nptr, wchar_t **
5484 restrict endptr, int base);
```

### **Description**

5485           The `wcstouq()` function shall convert the initial portion of the wide string `nptr` to  
 5486           unsigned long long int representation. It is identical to `wctoull()`.

### **Return Value**

5487           Refer to `wctoull()`.

### **Errors**

5488           Refer to `wctoull()`.

## wscanf

### Name

5489 wscanf — convert formatted input

### Description

5490 The `scanf()` family of functions shall behave as described in ISO POSIX (2003),  
 5491 except as noted below.

### Differences

5492 The `%s`, `%S` and `%[` conversion specifiers shall accept an option length modifier `a`,  
 5493 which shall cause a memory buffer to be allocated to hold the string converted. In  
 5494 such a case, the argument corresponding to the conversion specifier should be a  
 5495 reference to a pointer value that will receive a pointer to the allocated buffer. If there  
 5496 is insufficient memory to allocate a buffer, the function may set `errno` to `ENOMEM`  
 5497 and a conversion error results.

5498 **Note:** This directly conflicts with the ISO C (1999) usage of `%a` as a conversion specifier  
 5499 for hexadecimal float values. While this conversion specifier should be supported, a  
 5500 format specifier such as "`%aseconds`" will have a different meaning on an LSB  
 5501 conforming system.

## xdr\_u\_int

### Name

5502 `xdr_u_int` — library routines for external data representation

### Synopsis

5503 `int xdr_u_int(XDR * xdrs, unsigned int * up);`

### Description

5504 `xdr_u_int()` is a filter primitive that translates between C unsigned integers and  
 5505 their external representations.

### Return Value

5506 On success, 1 is returned. On error, 0 is returned.

## 13.6 Interfaces for libm

5507 Table 13-24 defines the library name and shared object name for the libm library

5508 **Table 13-24 libm Definition**

Library:	libm
SONAME:	See archLSB.

5509 The behavior of the interfaces in this library is specified by the following specifications:  
 5510

[ISO99] ISO C (1999)  
 [LSB] This Specification

5512 [SUSv2] SUSv2  
 [SUSv3] ISO POSIX (2003)

## 13.6.1 Math

### 13.6.1.1 Interfaces for Math

5514 An LSB conforming implementation shall provide the generic functions for Math  
 5515 specified in Table 13-25, with the full mandatory functionality as described in the  
 5516 referenced underlying specification.

5517 **Table 13-25 libm - Math Function Interfaces**

<code>__finite [ISOC99]</code>	<code>__finitef [ISOC99]</code>	<code>__finitel [ISOC99]</code>	<code>__fpclassify [LSB]</code>
<code>__fpclassifyf [LSB]</code>	<code>__signbit [ISOC99]</code>	<code>__signbitff [ISOC99]</code>	<code>acos [SUSv3]</code>
<code>acosf [SUSv3]</code>	<code>acosh [SUSv3]</code>	<code>acoshf [SUSv3]</code>	<code>acoshl [SUSv3]</code>
<code>acosl [SUSv3]</code>	<code>asin [SUSv3]</code>	<code>asinf [SUSv3]</code>	<code>asinh [SUSv3]</code>
<code>asinhf [SUSv3]</code>	<code>asinhl [SUSv3]</code>	<code>asinl [SUSv3]</code>	<code>atan [SUSv3]</code>
<code>atan2 [SUSv3]</code>	<code>atan2f [SUSv3]</code>	<code>atan2l [SUSv3]</code>	<code>atanf [SUSv3]</code>
<code>atanh [SUSv3]</code>	<code>atanhf [SUSv3]</code>	<code>atanhl [SUSv3]</code>	<code>atanl [SUSv3]</code>
<code>cabs [SUSv3]</code>	<code>cabsf [SUSv3]</code>	<code>cabsl [SUSv3]</code>	<code>cacos [SUSv3]</code>
<code>cacosf [SUSv3]</code>	<code>cacosh [SUSv3]</code>	<code>cacoshf [SUSv3]</code>	<code>cacoshl [SUSv3]</code>
<code>cacosl [SUSv3]</code>	<code>carg [SUSv3]</code>	<code>cargf [SUSv3]</code>	<code>cargl [SUSv3]</code>
<code>casin [SUSv3]</code>	<code>casinf [SUSv3]</code>	<code>casinh [SUSv3]</code>	<code>casinhf [SUSv3]</code>
<code>casinhl [SUSv3]</code>	<code>casinl [SUSv3]</code>	<code>catan [SUSv3]</code>	<code>catanf [SUSv3]</code>
<code>catanh [SUSv3]</code>	<code>catanhf [SUSv3]</code>	<code>catanhl [SUSv3]</code>	<code>catanl [SUSv3]</code>
<code>cbrt [SUSv3]</code>	<code>cbrtf [SUSv3]</code>	<code>cbrtl [SUSv3]</code>	<code>ccos [SUSv3]</code>
<code>ccosf [SUSv3]</code>	<code>ccosh [SUSv3]</code>	<code>ccoshf [SUSv3]</code>	<code>ccoshl [SUSv3]</code>
<code>ccosl [SUSv3]</code>	<code>ceil [SUSv3]</code>	<code>ceilf [SUSv3]</code>	<code>ceil [SUSv3]</code>
<code>cexp [SUSv3]</code>	<code>cexpf [SUSv3]</code>	<code>cexpl [SUSv3]</code>	<code>cimag [SUSv3]</code>
<code>cimagf [SUSv3]</code>	<code>cimagl [SUSv3]</code>	<code>clog [SUSv3]</code>	<code>clog10 [ISOC99]</code>
<code>clog10f [ISOC99]</code>	<code>clog10l [ISOC99]</code>	<code>clogf [SUSv3]</code>	<code>clogl [SUSv3]</code>
<code>conj [SUSv3]</code>	<code>conjf [SUSv3]</code>	<code>conjl [SUSv3]</code>	<code>copysign [SUSv3]</code>
<code>copysignf [SUSv3]</code>	<code>copysignl [SUSv3]</code>	<code>cos [SUSv3]</code>	<code>cosf [SUSv3]</code>
<code>cosh [SUSv3]</code>	<code>coshf [SUSv3]</code>	<code>coshl [SUSv3]</code>	<code>cosl [SUSv3]</code>
<code>cpow [SUSv3]</code>	<code>cpowf [SUSv3]</code>	<code>cpowl [SUSv3]</code>	<code>cproj [SUSv3]</code>
<code>cprojf [SUSv3]</code>	<code>cprojl [SUSv3]</code>	<code>creal [SUSv3]</code>	<code>crealf [SUSv3]</code>
<code>creall [SUSv3]</code>	<code>csin [SUSv3]</code>	<code>csinf [SUSv3]</code>	<code>csinh [SUSv3]</code>
<code>csinhf [SUSv3]</code>	<code>csinhl [SUSv3]</code>	<code>csinl [SUSv3]</code>	<code>csqrt [SUSv3]</code>

csqrft [SUSv3]	csqrfl [SUSv3]	ctan [SUSv3]	ctanf [SUSv3]
ctanh [SUSv3]	ctanhf [SUSv3]	ctanh1 [SUSv3]	ctanh1 [SUSv3]
dremf [ISOC99]	dreml [ISOC99]	erf [SUSv3]	erfc [SUSv3]
erfcf [SUSv3]	erfc1 [SUSv3]	erff [SUSv3]	erfl [SUSv3]
exp [SUSv3]	exp2 [SUSv3]	exp2f [SUSv3]	expf [SUSv3]
expl [SUSv3]	expm1 [SUSv3]	expm1f [SUSv3]	expm1l [SUSv3]
fabs [SUSv3]	fabsf [SUSv3]	fabsl [SUSv3]	fdim [SUSv3]
fdimf [SUSv3]	fdiml [SUSv3]	feclearexcept [SUSv3]	fegetenv [SUSv3]
fegetexceptflag [SUSv3]	fegetround [SUSv3]	feholdexcept [SUSv3]	feraiseexcept [SUSv3]
fesetenv [SUSv3]	fesetexceptflag [SUSv3]	fesetround [SUSv3]	fetestexcept [SUSv3]
feupdateenv [SUSv3]	finite [SUSv2]	finitef [ISOC99]	finitel [ISOC99]
floor [SUSv3]	floorf [SUSv3]	floorl [SUSv3]	fma [SUSv3]
fmaf [SUSv3]	fmal [SUSv3]	fmax [SUSv3]	fmaxf [SUSv3]
fmaxl [SUSv3]	fmin [SUSv3]	fminf [SUSv3]	fminl [SUSv3]
fmod [SUSv3]	fmodf [SUSv3]	fmodl [SUSv3]	frexp [SUSv3]
frexpf [SUSv3]	frexpl [SUSv3]	gamma [SUSv2]	gammaf [ISOC99]
gammal [ISOC99]	hypot [SUSv3]	hypotf [SUSv3]	hypotl [SUSv3]
ilogb [SUSv3]	ilogbf [SUSv3]	ilogbl [SUSv3]	j0 [SUSv3]
j0f [ISOC99]	j0l [ISOC99]	j1 [SUSv3]	j1f [ISOC99]
j1l [ISOC99]	jn [SUSv3]	jnf [ISOC99]	jnl [ISOC99]
ldexp [SUSv3]	ldexpf [SUSv3]	ldexpl [SUSv3]	lgamma [SUSv3]
lgamma_r [ISOC99]	lgammaf [SUSv3]	lgammaf_r [ISOC99]	lgammal [SUSv3]
lgammal_r [ISOC99]	llrint [SUSv3]	llrintf [SUSv3]	llrintl [SUSv3]
llround [SUSv3]	llroundf [SUSv3]	llroundl [SUSv3]	log [SUSv3]
log10 [SUSv3]	log10f [SUSv3]	log10l [SUSv3]	log1p [SUSv3]
log1pf [SUSv3]	log1pl [SUSv3]	log2 [SUSv3]	log2f [SUSv3]
log2l [SUSv3]	logb [SUSv3]	logbf [SUSv3]	logbl [SUSv3]
logf [SUSv3]	logl [SUSv3]	lrint [SUSv3]	lrintf [SUSv3]
lrintl [SUSv3]	lround [SUSv3]	lroundf [SUSv3]	lroundl [SUSv3]
matherr [ISOC99]	modf [SUSv3]	modff [SUSv3]	modfl [SUSv3]

	nan [SUSv3]	nanf [SUSv3]	nanl [SUSv3]	nearbyint [SUSv3]
	nearbyintf [SUSv3]	nearbyintl [SUSv3]	nextafter [SUSv3]	nextafterf [SUSv3]
	nextafterl [SUSv3]	nexttoward [SUSv3]	nexttowardf [SUSv3]	nexttowardl [SUSv3]
	pow [SUSv3]	pow10 [ISOC99]	pow10f [ISOC99]	pow10l [ISOC99]
	powf [SUSv3]	powl [SUSv3]	remainder [SUSv3]	remainderf [SUSv3]
	remainderl [SUSv3]	remquo [SUSv3]	remquof [SUSv3]	remquol [SUSv3]
	rint [SUSv3]	rintf [SUSv3]	rintl [SUSv3]	round [SUSv3]
	roundf [SUSv3]	roundl [SUSv3]	scalb [SUSv3]	scalbf [ISOC99]
	scalbl [ISOC99]	scalbln [SUSv3]	scalblnf [SUSv3]	scalblnl [SUSv3]
	scalbn [SUSv3]	scalbnf [SUSv3]	scalbnl [SUSv3]	significand [ISOC99]
	significandf [ISOC99]	significndl [ISOC99]	sin [SUSv3]	sincos [ISOC99]
	sincosf [ISOC99]	sincosl [ISOC99]	sinf [SUSv3]	sinh [SUSv3]
	sinhf [SUSv3]	sinhl [SUSv3]	sinl [SUSv3]	sqrt [SUSv3]
	sqrtf [SUSv3]	sqrtl [SUSv3]	tan [SUSv3]	tanf [SUSv3]
	tanh [SUSv3]	tanhf [SUSv3]	tanhl [SUSv3]	tanl [SUSv3]
	tgamma [SUSv3]	tgammaf [SUSv3]	tgammal [SUSv3]	trunc [SUSv3]
	truncf [SUSv3]	truncl [SUSv3]	y0 [SUSv3]	y0f [ISOC99]
	y0l [ISOC99]	y1 [SUSv3]	y1f [ISOC99]	y1l [ISOC99]
5518	yn [SUSv3]	ynf [ISOC99]	ynl [ISOC99]	

5519 An LSB conforming implementation shall provide the generic data interfaces for  
 5520 Math specified in Table 13-26, with the full mandatory functionality as described in  
 5521 the referenced underlying specification.

5522 **Table 13-26 libm - Math Data Interfaces**

5523 signgam [SUSv3]			
----------------------	--	--	--

## 13.7 Data Definitions for libm

5524 This section defines global identifiers and their values that are associated with  
 5525 interfaces contained in libm. These definitions are organized into groups that  
 5526 correspond to system headers. This convention is used as a convenience for the  
 5527 reader, and does not imply the existence of these headers, or their content. Where an  
 5528 interface is defined as requiring a particular system header file all of the data  
 5529 definitions for that system header file presented here shall be in effect.

5530 This section gives data definitions to promote binary application portability, not to  
 5531 repeat source interface definitions available elsewhere. System providers and  
 5532 application developers should use this ABI to supplement - not to replace - source  
 5533 interface definition specifications.

5534 This specification uses the ISO C (1999) C Language as the reference programming  
 5535 language, and data definitions are specified in ISO C format. The C language is used  
 5536 here as a convenient notation. Using a C language description of these data objects  
 5537 does not preclude their use by other programming languages.

### 13.7.1 complex.h

```
5538
5539 #define complex _Complex
5540
5541     extern double cabs(double complex);
5542     extern float cabsf(float complex);
5543     extern long double cabsl(long double complex);
5544     extern double complex cacos(double complex);
5545     extern float complex cacosf(float complex);
5546     extern double complex cacosh(double complex);
5547     extern float complex cacoshf(float complex);
5548     extern long double complex cacoshl(long double complex);
5549     extern long double complex cacosl(long double complex);
5550     extern double carg(double complex);
5551     extern float cargf(float complex);
5552     extern long double cargl(long double complex);
5553     extern double complex casin(double complex);
5554     extern float complex casinf(float complex);
5555     extern double complex casinh(double complex);
5556     extern float complex casinhf(float complex);
5557     extern long double complex casinhl(long double complex);
5558     extern long double complex casinl(long double complex);
5559     extern double complex catan(double complex);
5560     extern float complex catanf(float complex);
5561     extern double complex catanh(double complex);
5562     extern float complex catanhf(float complex);
5563     extern long double complex catanhl(long double complex);
5564     extern long double complex catanl(long double complex);
5565     extern double complex ccos(double complex);
5566     extern float complex ccosf(float complex);
5567     extern double complex ccosh(double complex);
5568     extern float complex ccoshf(float complex);
5569     extern long double complex ccoshl(long double complex);
5570     extern long double complex ccosl(long double complex);
5571     extern double complex cexp(double complex);
5572     extern float complex cexpf(float complex);
5573     extern long double complex cexpl(long double complex);
5574     extern double cimag(double complex);
5575     extern float cimagf(float complex);
5576     extern long double cimagsl(long double complex);
5577     extern double complex clog(double complex);
5578     extern float complex clog10f(float complex);
5579     extern long double complex clog10l(long double complex);
5580     extern float complex cloglf(float complex);
5581     extern long double complex clogl(long double complex);
5582     extern double complex conj(double complex);
5583     extern float complex conjf(float complex);
5584     extern long double complex conjl(long double complex);
5585     extern double complex cpow(double complex, double complex);
5586     extern float complex cpowf(float complex, float complex);
5587     extern long double complex cpowl(long double complex, long double
5588     complex);
```

```

5589     extern double complex cproj(double complex);
5590     extern float complex cprojf(float complex);
5591     extern long double complex cprojl(long double complex);
5592     extern double creal(double complex);
5593     extern float crealf(float complex);
5594     extern long double creall(long double complex);
5595     extern double complex csin(double complex);
5596     extern float complex csinf(float complex);
5597     extern double complex csinh(double complex);
5598     extern float complex csinhf(float complex);
5599     extern long double complex csinhl(long double complex);
5600     extern long double complex csinl(long double complex);
5601     extern double complex csqrt(double complex);
5602     extern float complex csqrtf(float complex);
5603     extern long double complex csqrtd(long double complex);
5604     extern double complex ctan(double complex);
5605     extern float complex ctanf(float complex);
5606     extern double complex tanh(double complex);
5607     extern float complex tanhf(float complex);
5608     extern long double complex tanhl(long double complex);
5609     extern long double complex tanl(long double complex);

```

### 13.7.2 fenv.h

```

5610
5611     extern int feclearexcept(int);
5612     extern int fegetenv(fenv_t *);
5613     extern int fegetexceptflag(fexcept_t *, int);
5614     extern int fegetround(void);
5615     extern int feholdexcept(fenv_t *);
5616     extern int feraiseexcept(int);
5617     extern int fesetenv(const fenv_t *);
5618     extern int fesetexceptflag(const fexcept_t *, int);
5619     extern int fesetround(int);
5620     extern int fetestexcept(int);
5621     extern int feupdateenv(const fenv_t *);

```

### 13.7.3 math.h

```

5622
5623     #define DOMAIN 1
5624     #define SING 2
5625
5626     struct exception {
5627         int type;
5628         char *name;
5629         double arg1;
5630         double arg2;
5631         double retval;
5632     };
5633
5634     #define FP_NAN 0
5635     #define FP_INFINITE 1
5636     #define FP_ZERO 2
5637     #define FP_SUBNORMAL 3
5638     #define FP_NORMAL 4
5639
5640     #define isnormal(x) (fpclassify (x) == FP_NORMAL)
5641     #define isfinite(x) \
5642         (sizeof (x) == sizeof (float) ? __finitef (x) : sizeof (x) == \
5643         sizeof (double) ? __finite (x) : __finitel (x))
5644     #define isinf(x) \
5645         (sizeof (x) == sizeof (float) ? __isinff (x) : sizeof (x) == sizeof \
5646         (double) ? __isinf (x) : __isinfl (x))

```

```

5647 #define isnan(x)          \
5648     (sizeof (x) == sizeof (float) ? __isnanf (x) : sizeof (x) == \
5649     sizeof (double) ? __isnan (x) : __isnanl (x))
5650
5651 #define HUGE_VAL          0x1.0p2047
5652 #define HUGE_VALF         0x1.0p255f
5653 #define HUGE_VALL        0x1.0p32767L
5654
5655 #define NAN      ((float)0x7fc00000UL)
5656 #define M_1_PI   0.31830988618379067154
5657 #define M_LOG10E 0.43429448190325182765
5658 #define M_2_PI   0.63661977236758134308
5659 #define M_LN2    0.69314718055994530942
5660 #define M_SQRT1_2 0.70710678118654752440
5661 #define M_PI_4   0.78539816339744830962
5662 #define M_2_SQRTPI 1.12837916709551257390
5663 #define M_SQRT2 1.41421356237309504880
5664 #define M_LOG2E 1.4426950408889634074
5665 #define M_PI_2   1.57079632679489661923
5666 #define M_LN10  2.30258509299404568402
5667 #define M_E     2.7182818284590452354
5668 #define M_PI    3.14159265358979323846
5669 #define INFINITY    HUGE_VALF
5670
5671 #define MATH_ERRNO    1
5672 #define MATH_ERREXCEPT 2
5673
5674 #define isunordered(u, v)          \
5675     (__extension__({ __typeof__(u) __u = (u); __typeof__(v) __v = \
5676     (v); fpclassify (__u) == FP_NAN || fpclassify (__v) == FP_NAN; }))
5677 #define islessgreater(x, y)        \
5678     (__extension__({ __typeof__(x) __x = (x); __typeof__(y) __y = \
5679     (y); !isunordered (__x, __y) && (__x < __y || __y < __x); }))
5680 #define isless(x,y)              \
5681     (__extension__({ __typeof__(x) __x = (x); __typeof__(y) __y = \
5682     (y); !isunordered (__x, __y) && __x < __y; }))
5683 #define islessequal(x, y)         \
5684     (__extension__({ __typeof__(x) __x = (x); __typeof__(y) __y = \
5685     (y); !isunordered (__x, __y) && __x <= __y; }))
5686 #define isgreater(x,y)           \
5687     (__extension__({ __typeof__(x) __x = (x); __typeof__(y) __y = \
5688     (y); !isunordered (__x, __y) && __x > __y; }))
5689 #define isgreaterequal(x,y)       \
5690     (__extension__({ __typeof__(x) __x = (x); __typeof__(y) __y = \
5691     (y); !isunordered (__x, __y) && __x >= __y; }))
5692
5693 extern int __finite(double);
5694 extern int __finitef(float);
5695 extern int __finitel(long double);
5696 extern int __isinf(double);
5697 extern int __isinff(float);
5698 extern int __isinfl(long double);
5699 extern int __isnan(double);
5700 extern int __isnanf(float);
5701 extern int __isnanl(long double);
5702 extern int __signbit(double);
5703 extern int __signbitf(float);
5704 extern int __fpclassify(double);
5705 extern int __fpclassifyf(float);
5706 extern int __fpclassifyl(long double);
5707 extern int signgam(void);
5708 extern double copysign(double, double);
5709 extern int finite(double);
5710 extern double frexp(double, int *);

```

```

5711     extern double ldexp(double, int);
5712     extern double modf(double, double *);
5713     extern double acos(double);
5714     extern double acosh(double);
5715     extern double asinh(double);
5716     extern double atanh(double);
5717     extern double asin(double);
5718     extern double atan(double);
5719     extern double atan2(double, double);
5720     extern double cbrt(double);
5721     extern double ceil(double);
5722     extern double cos(double);
5723     extern double cosh(double);
5724     extern double erf(double);
5725     extern double erfc(double);
5726     extern double exp(double);
5727     extern double expm1(double);
5728     extern double fabs(double);
5729     extern double floor(double);
5730     extern double fmod(double, double);
5731     extern double gamma(double);
5732     extern double hypot(double, double);
5733     extern int ilogb(double);
5734     extern double j0(double);
5735     extern double j1(double);
5736     extern double jn(int, double);
5737     extern double lgamma(double);
5738     extern double log(double);
5739     extern double log10(double);
5740     extern double log1p(double);
5741     extern double logb(double);
5742     extern double nextafter(double, double);
5743     extern double pow(double, double);
5744     extern double remainder(double, double);
5745     extern double rint(double);
5746     extern double scalb(double, double);
5747     extern double sin(double);
5748     extern double sinh(double);
5749     extern double sqrt(double);
5750     extern double tan(double);
5751     extern double tanh(double);
5752     extern double y0(double);
5753     extern double y1(double);
5754     extern double yn(int, double);
5755     extern float copysignf(float, float);
5756     extern long double copysignl(long double, long double);
5757     extern int finitef(float);
5758     extern int finitel(long double);
5759     extern float frexpf(float, int *);
5760     extern long double frexpl(long double, int *);
5761     extern float ldexpf(float, int);
5762     extern long double ldexpl(long double, int);
5763     extern float modff(float, float *);
5764     extern long double modfl(long double, long double *);
5765     extern double scalbln(double, long int);
5766     extern float scalblnf(float, long int);
5767     extern long double scalblnl(long double, long int);
5768     extern double scalbn(double, int);
5769     extern float scalbnf(float, int);
5770     extern long double scalbnl(long double, int);
5771     extern float acosf(float);
5772     extern float acoshf(float);
5773     extern long double acoshl(long double);
5774     extern long double acosl(long double);

```

```

5775 extern float asinf(float);
5776 extern float asinhf(float);
5777 extern long double asinhl(long double);
5778 extern long double asinl(long double);
5779 extern float atan2f(float, float);
5780 extern long double atan2l(long double, long double);
5781 extern float atanf(float);
5782 extern float atanhf(float);
5783 extern long double atanhl(long double);
5784 extern long double atanl(long double);
5785 extern float cbrtf(float);
5786 extern long double cbrtl(long double);
5787 extern float ceilf(float);
5788 extern long double ceill(long double);
5789 extern float cosf(float);
5790 extern float coshf(float);
5791 extern long double coshl(long double);
5792 extern long double cosl(long double);
5793 extern float dremf(float, float);
5794 extern long double dreml(long double, long double);
5795 extern float erfcf(float);
5796 extern long double erfcl(long double);
5797 extern float erff(float);
5798 extern long double erfl(long double);
5799 extern double exp2(double);
5800 extern float exp2f(float);
5801 extern long double exp2l(long double);
5802 extern float expf(float);
5803 extern long double expl(long double);
5804 extern float expmlf(float);
5805 extern long double expm1l(long double);
5806 extern float fabsf(float);
5807 extern long double fabsl(long double);
5808 extern double fdim(double, double);
5809 extern float fdimf(float, float);
5810 extern long double fdiml(long double, long double);
5811 extern float floorf(float);
5812 extern long double floorl(long double);
5813 extern double fma(double, double, double);
5814 extern float fmaf(float, float, float);
5815 extern long double fmål(long double, long double, long double);
5816 extern double fmax(double, double);
5817 extern float fmaxf(float, float);
5818 extern long double fmaxl(long double, long double);
5819 extern double fmin(double, double);
5820 extern float fminf(float, float);
5821 extern long double fminl(long double, long double);
5822 extern float fmodf(float, float);
5823 extern long double fmodl(long double, long double);
5824 extern float gammaf(float);
5825 extern long double gammal(long double);
5826 extern float hypotf(float, float);
5827 extern long double hypotl(long double, long double);
5828 extern int ilogbf(float);
5829 extern int ilogbl(long double);
5830 extern float j0f(float);
5831 extern long double j0l(long double);
5832 extern float j1f(float);
5833 extern long double j1l(long double);
5834 extern float jnf(int, float);
5835 extern long double jnl(int, long double);
5836 extern double lgamma_r(double, int *);
5837 extern float lgammaf(float);
5838 extern float lgammaf_r(float, int *);

```

```

5839     extern long double lgammal(long double);
5840     extern long double lgammal_r(long double, int *);
5841     extern long long int llrint(double);
5842     extern long long int llrintf(float);
5843     extern long long int llrintl(long double);
5844     extern long long int llround(double);
5845     extern long long int llroundf(float);
5846     extern long long int llroundl(long double);
5847     extern float log10f(float);
5848     extern long double log10l(long double);
5849     extern float log1pf(float);
5850     extern long double log1pl(long double);
5851     extern double log2(double);
5852     extern float log2f(float);
5853     extern long double log2l(long double);
5854     extern float logbf(float);
5855     extern long double logbl(long double);
5856     extern float logf(float);
5857     extern long double logl(long double);
5858     extern long int lrint(double);
5859     extern long int lrintf(float);
5860     extern long int lrintl(long double);
5861     extern long int lround(double);
5862     extern long int lroundf(float);
5863     extern long int lroundl(long double);
5864     extern int matherr(struct exception *);
5865     extern double nan(const char *);
5866     extern float nanf(const char *);
5867     extern long double nanl(const char *);
5868     extern double nearbyint(double);
5869     extern float nearbyintf(float);
5870     extern long double nearbyintl(long double);
5871     extern float nextafterf(float, float);
5872     extern long double nextafterl(long double, long double);
5873     extern double nexttoward(double, long double);
5874     extern float nexttowardf(float, long double);
5875     extern long double nexttowardl(long double, long double);
5876     extern double pow10(double);
5877     extern float pow10f(float);
5878     extern long double pow10l(long double);
5879     extern float powf(float, float);
5880     extern long double powl(long double, long double);
5881     extern float remainderf(float, float);
5882     extern long double remainderl(long double, long double);
5883     extern double remquo(double, double, int *);
5884     extern float remquof(float, float, int *);
5885     extern long double remquol(long double, long double, int *);
5886     extern float rintf(float);
5887     extern long double rintl(long double);
5888     extern double round(double);
5889     extern float roundf(float);
5890     extern long double roundl(long double);
5891     extern float scalbf(float, float);
5892     extern long double scalbl(long double, long double);
5893     extern double significand(double);
5894     extern float significandf(float);
5895     extern long double significndl(long double);
5896     extern void sincos(double *, double *);
5897     extern void sincosf(float, float *, float *);
5898     extern void sincosl(long double, long double *, long double *);
5899     extern float sinf(float);
5900     extern float sinhf(float);
5901     extern long double sinh1(long double);
5902     extern long double sinl(long double);

```

```

5903    extern float sqrtf(float);
5904    extern long double sqrtl(long double);
5905    extern float tanf(float);
5906    extern float tanhf(float);
5907    extern long double tanhl(long double);
5908    extern long double tanl(long double);
5909    extern double tgamma(double);
5910    extern float tgammaf(float);
5911    extern long double tgammal(long double);
5912    extern double trunc(double);
5913    extern float truncf(float);
5914    extern long double truncl(long double);
5915    extern float y0f(float);
5916    extern long double y0l(long double);
5917    extern float ylf(float);
5918    extern long double yll(long double);
5919    extern float ynf(int, float);
5920    extern long double ynl(int, long double);
5921    extern int __fpclassifyl(long double);
5922    extern int __fpclassifyl(long double);
5923    extern int __signbitl(long double);
5924    extern int __signbitl(long double);
5925    extern int __signbitl(long double);
5926    extern long double exp2l(long double);
5927    extern long double exp2l(long double);

```

## 13.8 Interface Definitions for libm

The interfaces defined on the following pages are included in libm and are defined by this specification. Unless otherwise noted, these interfaces shall be included in the source standard.

Other interfaces listed in Section 13.6 shall behave as described in the referenced base document.

### **\_\_fpclassify**

#### **Name**

5933       **\_\_fpclassify** – Classify real floating type

#### **Synopsis**

5934       int **\_\_fpclassify**(double *arg*);

#### **Description**

5935       **\_\_fpclassify**() has the same specification as **fpclassify**() in ISO POSIX (2003),  
 5936       except that the argument type for **\_\_fpclassify**() is known to be double.

5937       **\_\_fpclassify**() is not in the source standard; it is only in the binary standard.

## **\_\_fpclassifyf**

### **Name**

5938        \_\_fpclassifyf — Classify real floating type

### **Synopsis**

5939        int \_\_fpclassifyf(float arg);

### **Description**

5940        \_\_fpclassifyf() has the same specification as fpclassifyf() in ISO POSIX (2003),  
 5941        except that the argument type for \_\_fpclassifyf() is known to be float.

5942        \_\_fpclassifyf() is not in the source standard; it is only in the binary standard.

## **13.9 Interfaces for libpthread**

5943        Table 13-27 defines the library name and shared object name for the libpthread  
 5944        library

5945        **Table 13-27 libpthread Definition**

Library:	libpthread
SONAME:	libpthread.so.0

5947        The behavior of the interfaces in this library is specified by the following specifications:  
 5948

[LFS] Large File Support

[LSB] This Specification

[SUSv3] ISO POSIX (2003)

### **13.9.1 Realtime Threads**

#### **13.9.1.1 Interfaces for Realtime Threads**

5951        An LSB conforming implementation shall provide the generic functions for Realtime  
 5952        Threads specified in Table 13-28, with the full mandatory functionality as described  
 5953        in the referenced underlying specification.

5954        **Table 13-28 libpthread - Realtime Threads Function Interfaces**

pthread_attr_getinheritsched [SUSv3]	pthread_attr_getschedpolicy [SUSv3]	pthread_attr_gets cope [SUSv3]	pthread_attr_setinheritsched [SUSv3]
pthread_attr_setschedpolicy [SUSv3]	pthread_attr_setschedule [SUSv3]	pthread_getschedparam [SUSv3]	pthread_setschedparam [SUSv3]
pthread_setschedprio(GLIBC_2.3.4) [SUSv3]			

## 13.9.2 Advanced Realtime Threads

### 13.9.2.1 Interfaces for Advanced Realtime Threads

No external functions are defined for libpthread - Advanced Realtime Threads in this part of the specification. See also the relevant architecture specific supplement.

## 13.9.3 Posix Threads

### 13.9.3.1 Interfaces for Posix Threads

An LSB conforming implementation shall provide the generic functions for Posix Threads specified in Table 13-29, with the full mandatory functionality as described in the referenced underlying specification.

**Table 13-29 libpthread - Posix Threads Function Interfaces**

_pthread_cleanup_pop [LSB]	_pthread_cleanup_push [LSB]	pthread_attr_destroy [SUSv3]	pthread_attr_getdetachstate [SUSv3]
pthread_attr_getguardsize [SUSv3]	pthread_attr_getschedparam [SUSv3]	pthread_attr_getstack [SUSv3]	pthread_attr_getstackaddr [SUSv3]
pthread_attr_getstacksize [SUSv3]	pthread_attr_init [SUSv3]	pthread_attr_setdetachstate [SUSv3]	pthread_attr_setguardsize [SUSv3]
pthread_attr_setschedparam [SUSv3]	pthread_attr_setsstack [SUSv3]	pthread_attr_setsstackaddr [SUSv3]	pthread_attr_setsstacksize [SUSv3]
pthread_cancel [SUSv3]	pthread_cond_broadcast [SUSv3]	pthread_cond_destroy [SUSv3]	pthread_cond_init [SUSv3]
pthread_cond_signal [SUSv3]	pthread_cond_timedwait [SUSv3]	pthread_cond_wait [SUSv3]	pthread_condattr_destroy [SUSv3]
pthread_condattr_getpshared [SUSv3]	pthread_condattr_init [SUSv3]	pthread_condattr_setpshared [SUSv3]	pthread_create [SUSv3]
pthread_detach [SUSv3]	pthread_equal [SUSv3]	pthread_exit [SUSv3]	pthread_getconcurrency [SUSv3]
pthread_getspecific [SUSv3]	pthread_join [SUSv3]	pthread_key_create [SUSv3]	pthread_key_delete [SUSv3]
pthread_kill [SUSv3]	pthread_mutex_destory [SUSv3]	pthread_mutex_init [SUSv3]	pthread_mutex_lock [SUSv3]
pthread_mutex_trylock [SUSv3]	pthread_mutex_unlock [SUSv3]	pthread_mutexattr_destroy [SUSv3]	pthread_mutexattr_getpshared [SUSv3]
pthread_mutexattr_gettype [SUSv3]	pthread_mutexattr_init [SUSv3]	pthread_mutexattr_setpshared [SUSv3]	pthread_mutexattr_settype [SUSv3]
pthread_once	pthread_rwlock_destroy [SUSv3]	pthread_rwlock_init [SUSv3]	pthread_rwlock_rdlock [SUSv3]

5964

[SUSv3]	estroy [SUSv3]	nit [SUSv3]	dlock [SUSv3]
pthread_rwlock_timedlock [SUSv3]	pthread_rwlock_timedwrlock [SUSv3]	pthread_rwlock_tryrdlock [SUSv3]	pthread_rwlock_trywrlock [SUSv3]
pthread_rwlock_untime [SUSv3]	pthread_rwlock_wrlock [SUSv3]	pthread_rwlockattr_destroy [SUSv3]	pthread_rwlockattr_getshared [SUSv3]
pthread_rwlockattr_init [SUSv3]	pthread_rwlockattr_setshared [SUSv3]	pthread_self [SUSv3]	pthread_setcancelstate [SUSv3]
pthread_setcanceltype [SUSv3]	pthread_setconcurrency [SUSv3]	pthread_setspecific [SUSv3]	pthread_sigmask [SUSv3]
pthread_testcancel [SUSv3]	sem_close [SUSv3]	sem_destroy [SUSv3]	sem_getvalue [SUSv3]
sem_init [SUSv3]	sem_open [SUSv3]	sem_post [SUSv3]	sem_timedwait [SUSv3]
sem_trywait [SUSv3]	sem_unlink [SUSv3]	sem_wait [SUSv3]	

### 13.9.4 Thread aware versions of libc interfaces

5965

#### 13.9.4.1 Interfaces for Thread aware versions of libc interfaces

5966

5967

5968

An LSB conforming implementation shall provide the generic functions for Thread aware versions of libc interfaces specified in Table 13-30, with the full mandatory functionality as described in the referenced underlying specification.

5969

5970

**Table 13-30 libpthread - Thread aware versions of libc interfaces Function Interfaces**

5971

lseek64 [LFS]	open64 [LFS]	pread [SUSv3]	pread64 [LFS]
pwrite [SUSv3]	pwrite64 [LFS]		

## 13.10 Data Definitions for libpthread

5972

5973

5974

5975

5976

5977

This section defines global identifiers and their values that are associated with interfaces contained in libpthread. These definitions are organized into groups that correspond to system headers. This convention is used as a convenience for the reader, and does not imply the existence of these headers, or their content. Where an interface is defined as requiring a particular system header file all of the data definitions for that system header file presented here shall be in effect.

5978

5979

5980

5981

This section gives data definitions to promote binary application portability, not to repeat source interface definitions available elsewhere. System providers and application developers should use this ABI to supplement - not to replace - source interface definition specifications.

5982

5983

This specification uses the ISO C (1999) C Language as the reference programming language, and data definitions are specified in ISO C format. The C language is used

5984 here as a convenient notation. Using a C language description of these data objects  
 5985 does not preclude their use by other programming languages.

### 13.10.1 pthread.h

```

5986 #define PTHREAD_SCOPE_SYSTEM      0
5987 #define PTHREAD_MUTEX_DEFAULT     1
5988 #define PTHREAD_MUTEX_NORMAL      1
5989 #define PTHREAD_SCOPE_PROCESS     1
5990 #define PTHREAD_MUTEX_RECURSIVE   2
5991 #define PTHREAD_RWLOCK_DEFAULT_NP 2
5992 #define PTHREAD_MUTEX_ERRORCHECK   3
5993 #define PTHREAD_MUTEX_INITIALIZER \
5994     { 0, 0, 0, PTHREAD_MUTEX_NORMAL, __LOCK_INITIALIZER }
5995 #define PTHREAD_RWLOCK_INITIALIZER \
5996     { __LOCK_INITIALIZER, 0, NULL, NULL,
5997       NULL, PTHREAD_RWLOCK_DEFAULT_NP, \
5998       PTHREAD_PROCESS_PRIVATE }
5999 #define pthread_cleanup_push(routine,arg) \
6000     { struct __pthread_cleanup_buffer __buffer; \
6001       __pthread_cleanup_push(&__buffer,(routine),(arg)); \
6002     }
6003 #define pthread_cleanup_pop(execute) \
6004     __pthread_cleanup_pop(&__buffer,(execute)); \
6005 #define __LOCK_INITIALIZER { 0, 0 }
6006 #define PTHREAD_COND_INITIALIZER { __LOCK_INITIALIZER, 0 }
6007
6008 struct __pthread_cleanup_buffer {
6009     void (*__routine) (void *);
6010     void * __arg;
6011     int __canceltype;
6012     struct __pthread_cleanup_buffer *__prev;
6013 };
6014 typedef unsigned int pthread_key_t;
6015 typedef int pthread_once_t;
6016 typedef long long int __pthread_cond_align_t;
6017
6018 typedef unsigned long int pthread_t;
6019 struct __pthread_fastlock {
6020     long int __status;
6021     int __spinlock;
6022 };
6023
6024 typedef struct __pthread_descr_struct * __pthread_descr;
6025
6026 typedef struct {
6027     int __m_reserved;
6028     int __m_count;
6029     __pthread_descr __m_owner;
6030     int __m_kind;
6031     struct __pthread_fastlock __m_lock;
6032 } pthread_mutex_t;
6033 typedef struct {
6034     int __mutexkind;
6035 } pthread_mutexattr_t;
6036
6037 typedef struct {
6038     int __detachstate;
6039     int __schedpolicy;
6040     struct sched_param __schedparam;
6041     int __inheritsched;
6042     int __scope;
6043     size_t __guardsize;

```

```

6044         int __stackaddr_set;
6045         void *__stackaddr;
6046         unsigned long int __stacksize;
6047     } pthread_attr_t;
6048
6049     typedef struct {
6050         struct _pthread_fastlock __c_lock;
6051         _pthread_descr __c_waiting;
6052         char __padding[48 - sizeof(struct _pthread_fastlock) -
6053                         sizeof(_pthread_descr) -
6054                         sizeof(__pthread_cond_align_t)];
6055         __pthread_cond_align_t __align;
6056     } pthread_cond_t;
6057     typedef struct {
6058         int __dummy;
6059     } pthread_condattr_t;
6060
6061     typedef struct _pthread_rwlock_t {
6062         struct _pthread_fastlock __rw_lock;
6063         int __rw_readers;
6064         _pthread_descr __rw_writer;
6065         _pthread_descr __rw_read_waiting;
6066         _pthread_descr __rw_write_waiting;
6067         int __rw_kind;
6068         int __rw_pshared;
6069     } pthread_rwlock_t;
6070     typedef struct {
6071         int __lockkind;
6072         int __pshared;
6073     } pthread_rwlockattr_t;
6074
6075 #define PTHREAD_CREATE_JOINABLE 0
6076 #define PTHREAD_INHERIT_SCHED 0
6077 #define PTHREAD_ONCE_INIT 0
6078 #define PTHREAD_PROCESS_PRIVATE 0
6079 #define PTHREAD_CREATE_DETACHED 1
6080 #define PTHREAD_EXPLICIT_SCHED 1
6081 #define PTHREAD_PROCESS_SHARED 1
6082
6083 #define PTHREAD_CANCELED ((void*)-1)
6084 #define PTHREAD_CANCEL_DEFERRED 0
6085 #define PTHREAD_CANCEL_ENABLE 0
6086 #define PTHREAD_CANCEL_ASYNCHRONOUS 1
6087 #define PTHREAD_CANCEL_DISABLE 1
6088
6089 extern void _pthread_cleanup_pop(struct _pthread_cleanup_buffer *,
6090 int);
6091 extern void _pthread_cleanup_push(struct _pthread_cleanup_buffer *,
6092                                     void (*__routine) (void *)
6093                                     , void *);
6094 extern int pthread_attr_destroy(pthread_attr_t *);
6095 extern int pthread_attr_getdetachstate(const typedef struct {
6096                                         int __detachstate;
6097                                         int __schedpolicy;
6098                                         struct sched_param
6099                                         __schedparam;
6100                                         int __inheritsched;
6101                                         int __scope;
6102                                         size_t __guardsize;
6103                                         int __stackaddr_set;
6104                                         void *__stackaddr;
6105                                         unsigned long int __stacksize; }
6106                                         pthread_attr_t *, int *));
6107 extern int pthread_attr_getinheritsched(const typedef struct {

```

```

6108
6109
6110
6111     __schedparam;
6112
6113
6114
6115
6116
6117
6118     __stacksize; }
6119
6120     pthread_attr_t *, int *);
6121     extern int pthread_attr_getschedparam(const typedef struct {
6122         int __detachstate;
6123         int __schedpolicy;
6124         struct sched_param
6125             __schedparam;
6126
6127
6128
6129
6130
6131
6132     sched_param {
6133         int sched_priority; }
6134
6135
6136     * );
6137     extern int pthread_attr_getschedpolicy(const typedef struct {
6138         int __detachstate;
6139         int __schedpolicy;
6140         struct sched_param
6141             __schedparam;
6142
6143
6144
6145
6146
6147
6148     extern int pthread_attr_getscope(const typedef struct {
6149         int __detachstate;
6150         int __schedpolicy;
6151         struct sched_param __schedparam;
6152         int __inheritsched;
6153         int __scope;
6154         size_t __guardsize;
6155         int __stackaddr_set;
6156         void *__stackaddr;
6157         unsigned long int __stacksize; }
6158         pthread_attr_t *, int *);
6159     extern int pthread_attr_init(pthread_attr_t *);
6160     extern int pthread_attr_setdetachstate(pthread_attr_t *, int);
6161     extern int pthread_attr_setinheritsched(pthread_attr_t *, int);
6162     extern int pthread_attr_setschedparam(pthread_attr_t *, const struct
6163         sched_param {
6164             int sched_priority; }
6165
6166
6167     * );
6168     extern int pthread_attr_setschedpolicy(pthread_attr_t *, int);
6169     extern int pthread_attr_setscope(pthread_attr_t *, int);
6170     extern int pthread_cancel(typedef unsigned long int pthread_t);
6171     extern int pthread_cond_broadcast(pthread_cond_t *);
6172     extern int pthread_cond_destroy(pthread_cond_t *);

```

```

6172     extern int pthread_cond_init(pthread_cond_t *, const typedef struct {
6173             int __dummy; }
6174
6175             pthread_condattr_t *);
6176     extern int pthread_cond_signal(pthread_cond_t *);
6177     extern int pthread_cond_timedwait(pthread_cond_t *, pthread_mutex_t *,
6178     const struct timespec {
6179             time_t tv_sec; long int tv_nsec; }
6180
6181             * );
6182     extern int pthread_cond_wait(pthread_cond_t *, pthread_mutex_t *);
6183     extern int pthread_condattr_destroy(pthread_condattr_t *);
6184     extern int pthread_condattr_init(pthread_condattr_t *);
6185     extern int pthread_create(pthread_t *, const typedef struct {
6186             int __detachstate;
6187             int __schedpolicy;
6188             struct sched_param __schedparam;
6189             int __inheritsched;
6190             int __scope;
6191             size_t __guardsize;
6192             int __stackaddr_set;
6193             void *__stackaddr;
6194             unsigned long int __stacksize; }
6195             pthread_attr_t *,
6196             void *(*__start_routine) (void *p1)
6197             , void *); }
6198     extern int pthread_detach(pthread_t);
6199     extern int pthread_equal(pthread_t,
6200             pthread_t);
6201     extern void pthread_exit(void *); }
6202     extern int pthread_getschedparam(pthread_t,
6203             int *, struct sched_param {
6204             int sched_priority; }
6205
6206             * );
6207     extern void *pthread_getspecific(pthread_key_t); }
6208     extern int pthread_join(pthread_t, void **); }
6209     extern int pthread_key_create(pthread_key_t *, void (*destr_func) (void *))
6210             ); }
6211     extern int pthread_key_delete(pthread_key_t); }
6212     extern int pthread_mutex_destroy(pthread_mutex_t); }
6213     extern int pthread_mutex_init(pthread_mutex_t *, const typedef struct {
6214
6215             int __mutexkind; }
6216
6217             pthread_mutexattr_t *); }
6218     extern int pthread_mutex_lock(pthread_mutex_t *); }
6219     extern int pthread_mutex_trylock(pthread_mutex_t *); }
6220     extern int pthread_mutex_unlock(pthread_mutex_t *); }
6221     extern int pthread_mutexattr_destroy(pthread_mutexattr_t *); }
6222     extern int pthread_mutexattr_init(pthread_mutexattr_t *); }
6223     extern int pthread_once(pthread_once_t *, void (*init_routine) (void))
6224             ); }
6225     extern int pthread_rwlock_destroy(pthread_rwlock_t *); }
6226     extern int pthread_rwlock_init(pthread_rwlock_t *,
6227             pthread_rwlockattr_t *); }
6228     extern int pthread_rwlock_rdlock(pthread_rwlock_t *); }
6229     extern int pthread_rwlock_tryrdlock(pthread_rwlock_t *); }
6230     extern int pthread_rwlock_trywrlock(pthread_rwlock_t *); }
6231     extern int pthread_rwlock_unlock(pthread_rwlock_t *); }
6232     extern int pthread_rwlock_wrlock(pthread_rwlock_t *); }
6233     extern int pthread_rwlockattr_destroy(pthread_rwlockattr_t *); }
6234     extern int pthread_rwlockattr_getpshared(const typedef struct {
```

```

6236                               int __lockkind; int
6237   __pshared; }
6238
6239   );
6240   extern int pthread_rwlockattr_init(pthread_rwlockattr_t *);
6241   extern int pthread_rwlockattr_setpshared(pthread_rwlockattr_t *, int);
6242   extern typedef unsigned long int pthread_t pthread_self(void);
6243   extern int pthread_setcancelstate(int, int *);
6244   extern int pthread_setcanceltype(int, int *);
6245   extern int pthread_setschedparam(pthread_attr_t, int,
6246   const struct sched_param {
6247           int sched_priority;
6248
6249   });
6250   extern int pthread_setspecific(pthread_key_t, const void *);
6251
6252   extern void pthread_testcancel(void);
6253   extern int pthread_attr_getguardsize(const pthread_attr_t {
6254           int __detachstate;
6255           int __schedpolicy;
6256           struct sched_param __schedparam;
6257           int __inheritsched;
6258           int __scope;
6259           size_t __guardsize;
6260           int __stackaddr_set;
6261           void * __stackaddr;
6262           unsigned long int __stacksize;
6263           pthread_attr_t *, size_t *);
6264   extern int pthread_attr_setguardsize(pthread_attr_t *,
6265           typedef unsigned long int
6266           size_t);
6267   extern int pthread_attr_setstackaddr(pthread_attr_t *, void *);
6268   extern int pthread_attr_getstackaddr(const pthread_attr_t {
6269           int __detachstate;
6270           int __schedpolicy;
6271           struct sched_param __schedparam;
6272           int __inheritsched;
6273           int __scope;
6274           size_t __guardsize;
6275           int __stackaddr_set;
6276           void * __stackaddr;
6277           unsigned long int __stacksize;
6278           pthread_attr_t *, void **);
6279   extern int pthread_attr_setstacksize(pthread_attr_t *,
6280           typedef unsigned long int
6281           size_t);
6282   extern int pthread_attr_getstacksize(const pthread_attr_t {
6283           int __detachstate;
6284           int __schedpolicy;
6285           struct sched_param __schedparam;
6286           int __inheritsched;
6287           int __scope;
6288           size_t __guardsize;
6289           int __stackaddr_set;
6290           void * __stackaddr;
6291           unsigned long int __stacksize;
6292           pthread_attr_t *, size_t *);
6293   extern int pthread_mutexattr_gettype(const pthread_mutexattr_t {
6294           int __mutexkind;
6295           pthread_mutexattr_t *, int *);
6296   extern int pthread_mutexattr_settype(pthread_mutexattr_t *, int);
6297   extern int pthread_getconcurrency(void);
6298   extern int pthread_setconcurrency(int);
6299   extern int pthread_attr_getstack(const pthread_attr_t {

```

```

6300                               int __detachstate;
6301                               int __schedpolicy;
6302                               struct sched_param __schedparam;
6303                               int __inheritsched;
6304                               int __scope;
6305                               size_t __guardsize;
6306                               int __stackaddr_set;
6307                               void * __stackaddr;
6308                               unsigned long int __stacksize; }
6309                               pthread_attr_t *, void **, size_t *);
6310 extern int pthread_attr_setstack(pthread_attr_t *, void *,
6311                                     typedef unsigned long int size_t);
6312 extern int pthread_condattr_getpshared(const typedef struct {
6313                                         int __dummy; }
6314                                         pthread_condattr_t *, int *));
6315 extern int pthread_condattr_setpshared(pthread_condattr_t *, int);
6316 extern int pthread_mutexattr_getpshared(const typedef struct {
6317                                         int __mutexkind; }
6318                                         pthread_mutexattr_t *, int *));
6319 extern int pthread_mutexattr_setpshared(pthread_mutexattr_t *, int);
6320 extern int pthread_rwlock_timedrdlock(pthread_rwlock_t *, const struct
6321 timespec {                                               time_t tv_sec; long int
6322                                         tv_nsec; }
6323                                         * );
6324
6325 extern int pthread_rwlock_timedwrlock(pthread_rwlock_t *, const struct
6326 timespec {                                               time_t tv_sec; long int
6327                                         tv_nsec; }
6328                                         * );
6329
6330 extern int __register_atfork(void (*prepare) (void)
6331                               , void (*parent) (void)
6332                               , void (*child) (void)
6333                               , void * );
6334 extern int pthread_setschedprio(pthread_attr_t,
6335 int);
```

## 13.10.2 semaphore.h

```

6338
6339     typedef struct {
6340         struct _pthread_fastlock __sem_lock;
6341         int __sem_value;
6342         _pthread_descr __sem_waiting;
6343     } sem_t;
6344
6345 #define SEM_FAILED      ((sem_t*)0)
6346
6347 #define SEM_VALUE_MAX    ((int)((~0u)>>1))
6348
6349 extern int sem_close(sem_t *);
6350 extern int sem_destroy(sem_t *);
6351 extern int sem_getvalue(sem_t *, int *);
6352 extern int sem_init(sem_t *, int, unsigned int);
6353 extern sem_t *sem_open(const char *, int, ...);
6354 extern int sem_post(sem_t *);
6355 extern int sem_trywait(sem_t *);
6356 extern int sem_unlink(const char *);
6357 extern int sem_wait(sem_t *);
6358 extern int sem_timedwait(sem_t *, const struct timespec *);
```

## 13.11 Interface Definitions for libpthread

6359 The interfaces defined on the following pages are included in libpthread and are  
 6360 defined by this specification. Unless otherwise noted, these interfaces shall be  
 6361 included in the source standard.

6362 Other interfaces listed in Section 13.9 shall behave as described in the referenced  
 6363 base document.

### \_pthread\_cleanup\_pop

#### Name

6364 `_pthread_cleanup_pop` — establish cancellation handlers

#### Synopsis

```
6365 #include <pthread.h>
6366 void _pthread_cleanup_pop(struct _pthread_cleanup_buffer *, int);
```

#### Description

6367 The `_pthread_cleanup_pop()` function provides an implementation of the  
 6368 `pthread_cleanup_pop()` macro described in *ISO POSIX (2003)*.

6369 The `_pthread_cleanup_pop()` function is not in the source standard; it is only in  
 6370 the binary standard.

### \_pthread\_cleanup\_push

#### Name

6371 `_pthread_cleanup_push` — establish cancellation handlers

#### Synopsis

```
6372 #include <pthread.h>
6373 void _pthread_cleanup_push(struct _pthread_cleanup_buffer *, void (*)
6374 (void *), void *);
```

#### Description

6375 The `_pthread_cleanup_push()` function provides an implementation of the  
 6376 `pthread_cleanup_push()` macro described in *ISO POSIX (2003)*.

6377 The `_pthread_cleanup_push()` function is not in the source standard; it is only in  
 6378 the binary standard.

## 13.12 Interfaces for libgcc\_s

6379 Table 13-31 defines the library name and shared object name for the libgcc\_s library

#### Table 13-31 libgcc\_s Definition

Library:	libgcc_s
SONAME:	libgcc_s.so.1

### 13.12.1 Unwind Library

#### 13.12.1.1 Interfaces for Unwind Library

No external functions are defined for libgcc\_s - Unwind Library in this part of the specification. See also the relevant architecture specific supplement.

## 13.13 Data Definitions for libgcc\_s

This section defines global identifiers and their values that are associated with interfaces contained in libgcc\_s. These definitions are organized into groups that correspond to system headers. This convention is used as a convenience for the reader, and does not imply the existence of these headers, or their content. Where an interface is defined as requiring a particular system header file all of the data definitions for that system header file presented here shall be in effect.

This section gives data definitions to promote binary application portability, not to repeat source interface definitions available elsewhere. System providers and application developers should use this ABI to supplement - not to replace - source interface definition specifications.

This specification uses the ISO C (1999) C Language as the reference programming language, and data definitions are specified in ISO C format. The C language is used here as a convenient notation. Using a C language description of these data objects does not preclude their use by other programming languages.

### 13.13.1 unwind.h

```

6399
6400     struct _Unwind_Context;
6401
6402     typedef void *_Unwind_Ptr;
6403     typedef unsigned int _Unwind_Word;
6404
6405     typedef enum {
6406         _URC_NO_REASON, _URC_FOREIGN_EXCEPTION_CAUGHT =
6407             1, _URC_FATAL_PHASE2_ERROR = 2, _URC_FATAL_PHASE1_ERROR =
6408             3, _URC_NORMAL_STOP = 4, _URC_END_OF_STACK =
6409             5, _URC_HANDLER_FOUND = 6, _URC_INSTALL_CONTEXT =
6410             7, _URC_CONTINUE_UNWIND = 8
6411     } _Unwind_Reason_Code;
6412
6413     struct _Unwind_Exception {
6414         u_int64_t exception_class;
6415         _Unwind_Exception_Cleanup_Fn exception_cleanup;
6416         u_int64_t private_1;
6417         u_int64_t private_2;
6418     };
6419
6420     #define _UA_SEARCH_PHASE          1
6421     #define _UA_END_OF_STACK          16
6422     #define _UA_CLEANUP_PHASE         2
6423     #define _UA_HANDLER_FRAME          4
6424     #define _UA_FORCE_UNWIND           8
6425
6426     extern void _Unwind_DeleteException(struct _Unwind_Exception *);
6427     extern fde *_Unwind_Find_FDE(void *, struct dwarf_eh_base *);
6428     extern void _Unwind_DeleteException(struct _Unwind_Exception *);
6429     extern _Unwind_Ptr _Unwind_ForcedUnwind(struct _Unwind_Exception *,
6430                                         _Unwind_Stop_Fn, void *);
6431     extern _Unwind_Word _Unwind_GetGR(struct _Unwind_Context *, int);

```

```

6432 extern _Unwind_Ptr _Unwind_GetIP(struct _Unwind_Context *);
6433 extern _Unwind_Ptr _Unwind_GetLanguageSpecificData(struct
6434 _Unwind_Context
6435 * );
6436 extern _Unwind_Ptr _Unwind_GetRegionStart(struct _Unwind_Context *);
6437 extern _Unwind_Reason_Code _Unwind_RaiseException(struct
6438 _Unwind_Exception
6439 * );
6440 extern void _Unwind_Resume(struct _Unwind_Exception *);
6441 extern void _Unwind_SetGR(struct _Unwind_Context *, int, u_int64_t);
6442 extern void _Unwind_SetIP(struct _Unwind_Context *, _Unwind_Ptr);
6443 extern void _Unwind_DeleteException(struct _Unwind_Exception *);
6444 extern fde *_Unwind_Find_FDE(void *, struct dwarf_eh_base *);
6445 extern _Unwind_Ptr _Unwind_ForcedUnwind(struct _Unwind_Exception *,
6446 _Unwind_Stop_Fn, void *);
6447 extern _Unwind_Ptr _Unwind_GetDataRelBase(struct _Unwind_Context *);
6448 extern _Unwind_Word _Unwind_GetGR(struct _Unwind_Context *, int);
6449 extern _Unwind_Ptr _Unwind_GetIP(struct _Unwind_Context *);
6450 extern _Unwind_Ptr _Unwind_GetLanguageSpecificData(struct
6451 _Unwind_Context
6452 * );
6453 extern _Unwind_Ptr _Unwind_GetRegionStart(struct _Unwind_Context *);
6454 extern _Unwind_Ptr _Unwind_GetTextRelBase(struct _Unwind_Context *);
6455 extern _Unwind_Reason_Code _Unwind_RaiseException(struct
6456 _Unwind_Exception
6457 * );
6458 extern void _Unwind_Resume(struct _Unwind_Exception *);
6459 extern void _Unwind_SetGR(struct _Unwind_Context *, int, u_int64_t);
6460 extern void _Unwind_SetIP(struct _Unwind_Context *, _Unwind_Ptr);
6461 extern _Unwind_Ptr _Unwind_ForcedUnwind(struct _Unwind_Exception *,
6462 _Unwind_Stop_Fn, void *);
6463 extern _Unwind_Ptr _Unwind_GetDataRelBase(struct _Unwind_Context *);
6464 extern _Unwind_Word _Unwind_GetGR(struct _Unwind_Context *, int);
6465 extern _Unwind_Ptr _Unwind_GetIP(struct _Unwind_Context *);
6466 extern _Unwind_Ptr _Unwind_GetLanguageSpecificData(struct
6467 _Unwind_Context
6468 * );
6469 extern _Unwind_Ptr _Unwind_GetRegionStart(struct _Unwind_Context *);
6470 extern _Unwind_Ptr _Unwind_GetTextRelBase(struct _Unwind_Context *);
6471 extern _Unwind_Reason_Code _Unwind_RaiseException(struct
6472 _Unwind_Exception
6473 * );
6474 extern void _Unwind_Resume(struct _Unwind_Exception *);
6475 extern void _Unwind_SetGR(struct _Unwind_Context *, int, u_int64_t);
6476 extern void _Unwind_SetIP(struct _Unwind_Context *, _Unwind_Ptr);
6477 extern void _Unwind_DeleteException(struct _Unwind_Exception *);
6478 extern fde *_Unwind_Find_FDE(void *, struct dwarf_eh_base *);
6479 extern _Unwind_Ptr _Unwind_ForcedUnwind(struct _Unwind_Exception *,
6480 _Unwind_Stop_Fn, void *);
6481 extern _Unwind_Ptr _Unwind_GetDataRelBase(struct _Unwind_Context *);
6482 extern _Unwind_Word _Unwind_GetGR(struct _Unwind_Context *, int);
6483 extern _Unwind_Ptr _Unwind_GetIP(struct _Unwind_Context *);
6484 extern _Unwind_Ptr _Unwind_GetLanguageSpecificData(struct
6485 _Unwind_Context
6486 * );
6487 extern _Unwind_Ptr _Unwind_GetRegionStart(struct _Unwind_Context *);
6488 extern _Unwind_Ptr _Unwind_GetTextRelBase(struct _Unwind_Context *);
6489 extern _Unwind_Reason_Code _Unwind_RaiseException(struct
6490 _Unwind_Exception
6491 * );
6492 extern void _Unwind_Resume(struct _Unwind_Exception *);
6493 extern void _Unwind_SetGR(struct _Unwind_Context *, int, u_int64_t);
6494 extern void _Unwind_SetIP(struct _Unwind_Context *, _Unwind_Ptr);
6495 extern void _Unwind_DeleteException(struct _Unwind_Exception *);

```

```

6496     extern fde *_Unwind_Find_FDE(void *, struct dwarf_eh_base *);
6497     extern _Unwind_Ptr _Unwind_ForcedUnwind(struct _Unwind_Exception *,
6498                                              _Unwind_Stop_Fn, void *);
6499     extern _Unwind_Ptr _Unwind_GetDataRelBase(struct _Unwind_Context *,
6500                                              _Unwind_Stop_Fn, void *);
6501     extern _Unwind_Word _Unwind_GetGR(struct _Unwind_Context *, int);
6502     extern _Unwind_Ptr _Unwind_GetIP(struct _Unwind_Context *);
6503     extern _Unwind_Ptr _Unwind_GetLanguageSpecificData(struct
6504                                              _Unwind_Context
6505                                              * );
6506     extern _Unwind_Ptr _Unwind_GetRegionStart(struct _Unwind_Context *);
6507     extern _Unwind_Ptr _Unwind_GetTextRelBase(struct _Unwind_Context *);
6508     extern _Unwind_Reason_Code _Unwind_RaiseException(struct
6509                                              _Unwind_Exception
6510                                              * );
6511     extern void _Unwind_Resume(struct _Unwind_Exception *);
6512     extern void _Unwind_SetGR(struct _Unwind_Context *, int, u_int64_t);
6513     extern void _Unwind_SetIP(struct _Unwind_Context *, _Unwind_Ptr);
6514     extern void _Unwind_DeleteException(struct _Unwind_Exception *);
6515     extern fde *_Unwind_Find_FDE(void *, struct dwarf_eh_base *);
6516     extern _Unwind_Ptr _Unwind_ForcedUnwind(struct _Unwind_Exception *,
6517                                              _Unwind_Stop_Fn, void *);
6518     extern _Unwind_Ptr _Unwind_GetDataRelBase(struct _Unwind_Context *);
6519     extern _Unwind_Word _Unwind_GetGR(struct _Unwind_Context *, int);
6520     extern _Unwind_Ptr _Unwind_GetIP(struct _Unwind_Context *);
6521     extern _Unwind_Ptr _Unwind_GetLanguageSpecificData(void);
6522     extern _Unwind_Ptr _Unwind_GetRegionStart(struct _Unwind_Context *);
6523     extern _Unwind_Ptr _Unwind_GetTextRelBase(struct _Unwind_Context *);
6524     extern _Unwind_Reason_Code _Unwind_RaiseException(struct
6525                                              _Unwind_Exception
6526                                              * );
6527     extern void _Unwind_Resume(struct _Unwind_Exception *);
6528     extern void _Unwind_SetGR(struct _Unwind_Context *, int, u_int64_t);
6529     extern void _Unwind_SetIP(struct _Unwind_Context *, _Unwind_Ptr);
6530     extern void _Unwind_DeleteException(struct _Unwind_Exception *);
6531     extern fde *_Unwind_Find_FDE(void *, struct dwarf_eh_base *);
6532     extern _Unwind_Ptr _Unwind_ForcedUnwind(struct _Unwind_Exception *,
6533                                              _Unwind_Stop_Fn, void *);
6534     extern _Unwind_Ptr _Unwind_GetDataRelBase(struct _Unwind_Context *);
6535     extern _Unwind_Word _Unwind_GetGR(struct _Unwind_Context *, int);
6536     extern _Unwind_Ptr _Unwind_GetIP(struct _Unwind_Context *);
6537     extern _Unwind_Ptr _Unwind_GetLanguageSpecificData(void);
6538     extern _Unwind_Ptr _Unwind_GetRegionStart(struct _Unwind_Context *);
6539     extern _Unwind_Ptr _Unwind_GetTextRelBase(struct _Unwind_Context *);
6540     extern _Unwind_Reason_Code _Unwind_RaiseException(struct
6541                                              _Unwind_Exception
6542                                              * );
6543     extern void _Unwind_Resume(struct _Unwind_Exception *);
6544     extern void _Unwind_SetGR(struct _Unwind_Context *, int, u_int64_t);
6545     extern void _Unwind_SetIP(struct _Unwind_Context *, _Unwind_Ptr);
6546     extern _Unwind_Reason_Code _Unwind_Backtrace(_Unwind_Trace_Fn, void
6547                                              * );
6548     extern _Unwind_Reason_Code _Unwind_Backtrace(_Unwind_Trace_Fn, void
6549                                              * );
6550     extern _Unwind_Reason_Code _Unwind_Backtrace(_Unwind_Trace_Fn, void
6551                                              * );
6552     extern _Unwind_Reason_Code _Unwind_Backtrace(_Unwind_Trace_Fn, void
6553                                              * );
6554     extern _Unwind_Reason_Code _Unwind_Backtrace(_Unwind_Trace_Fn, void
6555                                              * );
6556     extern _Unwind_Reason_Code _Unwind_Backtrace(_Unwind_Trace_Fn, void
6557                                              * );
6558     extern _Unwind_Reason_Code _Unwind_Backtrace(_Unwind_Trace_Fn, void
6559                                              * );
6559     extern _Unwind_Reason_Code _Unwind_GetCFA(struct _Unwind_Context *);

```

```

6560 extern _Unwind_Reason_Code _Unwind_GetCFA(struct _Unwind_Context * );
6561 extern _Unwind_Reason_Code _Unwind_GetCFA(struct _Unwind_Context * );
6562 extern _Unwind_Reason_Code _Unwind_GetCFA(struct _Unwind_Context * );
6563 extern _Unwind_Reason_Code _Unwind_GetCFA(struct _Unwind_Context * );
6564 extern _Unwind_Reason_Code _Unwind_GetCFA(struct _Unwind_Context * );
6565 extern _Unwind_Reason_Code _Unwind_GetCFA(struct _Unwind_Context * );
6566 extern _Unwind_Reason_Code _Unwind_Resume_or_Rethrow(struct
6567
6568 _Unwind_Exception * );
6569 extern _Unwind_Reason_Code _Unwind_Resume_or_Rethrow(struct
6570
6571 _Unwind_Exception * );
6572 extern _Unwind_Reason_Code _Unwind_Resume_or_Rethrow(struct
6573
6574 _Unwind_Exception * );
6575 extern _Unwind_Reason_Code _Unwind_Resume_or_Rethrow(struct
6576
6577 _Unwind_Exception * );
6578 extern _Unwind_Reason_Code _Unwind_Resume_or_Rethrow(struct
6579
6580 _Unwind_Exception * );
6581 extern _Unwind_Reason_Code _Unwind_Resume_or_Rethrow(struct
6582
6583 _Unwind_Exception * );
6584 extern _Unwind_Reason_Code _Unwind_Resume_or_Rethrow(struct
6585
6586 _Unwind_Exception * );
6587 extern void *_Unwind_FindEnclosingFunction(void * );
6588 extern void *_Unwind_FindEnclosingFunction(void * );
6589 extern void *_Unwind_FindEnclosingFunction(void * );
6590 extern void *_Unwind_FindEnclosingFunction(void * );
6591 extern void *_Unwind_FindEnclosingFunction(void * );
6592 extern void *_Unwind_FindEnclosingFunction(void * );
6593 extern void *_Unwind_FindEnclosingFunction(void * );
6594 extern _Unwind_Word _Unwind_GetBSP(struct _Unwind_Context * );

```

## 13.14 Interfaces for libdl

6595 Table 13-32 defines the library name and shared object name for the libdl library

6596 **Table 13-32 libdl Definition**

6597 Library:	libdl
SONAME:	libdl.so.2

6598 The behavior of the interfaces in this library is specified by the following specifications:  
 6599

[LSB] This Specification  
 6600 [SUSv3] ISO POSIX (2003)

### 13.14.1 Dynamic Loader

#### 13.14.1.1 Interfaces for Dynamic Loader

6601 An LSB conforming implementation shall provide the generic functions for Dynamic  
 6602 Loader specified in Table 13-33, with the full mandatory functionality as described  
 6603 in the referenced underlying specification.  
 6604

6605      **Table 13-33 libdl - Dynamic Loader Function Interfaces**

6606	dladdr [LSB]	dlclose [SUSv3]	dlerror [SUSv3]	dlopen [LSB]
	dlsym [LSB]			

## 13.15 Data Definitions for libdl

6607      This section defines global identifiers and their values that are associated with  
 6608      interfaces contained in libdl. These definitions are organized into groups that  
 6609      correspond to system headers. This convention is used as a convenience for the  
 6610      reader, and does not imply the existence of these headers, or their content. Where an  
 6611      interface is defined as requiring a particular system header file all of the data  
 6612      definitions for that system header file presented here shall be in effect.

6613      This section gives data definitions to promote binary application portability, not to  
 6614      repeat source interface definitions available elsewhere. System providers and  
 6615      application developers should use this ABI to supplement - not to replace - source  
 6616      interface definition specifications.

6617      This specification uses the ISO C (1999) C Language as the reference programming  
 6618      language, and data definitions are specified in ISO C format. The C language is used  
 6619      here as a convenient notation. Using a C language description of these data objects  
 6620      does not preclude their use by other programming languages.

### 13.15.1 dlfcn.h

```
6621
6622     #define RTLD_NEXT      ((void *) -11)
6623     #define RTLD_LOCAL      0
6624     #define RTLD_LAZY       0x00001
6625     #define RTLD_NOW        0x00002
6626     #define RTLD_GLOBAL      0x00100
6627
6628     typedef struct {
6629         char *dli_fname;
6630         void *dli_fbase;
6631         char *dli_sname;
6632         void *dli_saddr;
6633     } Dl_info;
6634     extern int dladdr(const void *, Dl_info *);
6635     extern int dlclose(void *);
6636     extern char *dlerror(void);
6637     extern void *dlopen(char *, int);
6638     extern void *dlsym(void *, char *);
```

## 13.16 Interface Definitions for libdl

6639      The interfaces defined on the following pages are included in libdl and are defined  
 6640      by this specification. Unless otherwise noted, these interfaces shall be included in the  
 6641      source standard.

6642      Other interfaces listed in Section 13.14 shall behave as described in the referenced  
 6643      base document.

## dladdr

### Name

6644        `dladdr` — find the shared object containing a given address

### Synopsis

```
6645 #include <dlfcn.h>
6646
6647 typedef struct {
6648     const char    *dli_fname;
6649     void           *dli_fbase;
6650     const char    *dli_sname;
6651     void           *dli_saddr;
```

```
6652     } Dl_info;
6653
6654     int dladdr(const void * addr, Dl_info * dlip);
```

## Description

6654 The `dladdr()` function shall query the dynamic linker for information about the  
 6655 shared object containing the address `addr`. The information shall be returned in the  
 6656 user supplied data structure referenced by `dlip`.

6657 The structure shall contain at least the following members:

```
6658     dli_fname
6659         The pathname of the shared object containing the address
6660
6661     dli_fbase
6662         The base address at which the shared object is mapped into the address space of
6663
6664     dli_sname
6665         The name of the nearest runtime symbol with value less than or equal to addr.
6666         Where possible, the symbol name shall be returned as it would appear in C
6667         source code.
6668         If no symbol with a suitable value is found, both this field and dli_saddr shall
6669         be set to NULL.
6670
6671     dli_saddr
6672         The address of the symbol returned in dli_sname. This address has type
6673         "pointer to type", where type is the type of the symbol dli_sname.
```

6674 **Example:** If the symbol in `dli_sname` is a function, then the type of `dli_saddr` is of type  
 6675 "pointer to function".

6676 The behavior of `dladdr()` is only specified in dynamically linked programs.

## Return Value

6677 On success, `dladdr()` shall return non-zero, and the structure referenced by `dli`  
 6678 shall be filled in as described. Otherwise, `dladdr()` shall return zero, and the cause  
 6679 of the error can be fetched with `dlerror()`.

## Errors

6680 See `dlerror()`.

## Environment

```
6681     LD_LIBRARY_PATH
6682         directory search-path for object files
```

## dlopen

### Name

6681 dlopen — open dynamic object

### Synopsis

```
6682 #include <dlfcn.h>
6683 void * dlopen(const char * filename, int flag);
```

### Description

6684 The `dlopen()` function shall behave as specified in ISO POSIX (2003), but with  
6685 additional behaviors listed below.

6686 If the file argument does not contain a slash character, then the system shall look for  
6687 a library of that name in at least the following directories, and use the first one which  
6688 is found:

- 6689 • The directories specified by the `DT_RPATH` dynamic entry.
- 6690 • The directories specified in the `LD_LIBRARY_PATH` environment variable (which is  
6691 a colon separated list of pathnames). This step shall be skipped for setuid and  
6692 setgid executables.
- 6693 • A set of directories sufficient to contain the libraries specified in this standard.

6694 **Note:** Traditionally, `/lib` and `/usr/lib`. This case would also cover cases in which the  
6695 system used the mechanism of `/etc/ld.so.conf` and `/etc/ld.so.cache` to provide  
6696 access.

6697 Example: An application which is not linked against `libm` may choose to `dlopen libm`.

## dlsym

### Name

6698 `dlsym` — obtain the address of a symbol from a `dlopen` object

### Description

6699 `dlsym()` is as specified in the ISO POSIX (2003), but with differences as listed below.

#### The special purpose value for handle RTLD\_NEXT

6700 The value `RTLD_NEXT`, which is reserved for future use shall be available, with the  
6701 behavior as described in ISO POSIX (2003).

## 13.17 Interfaces for librt

6703 Table 13-34 defines the library name and shared object name for the `librt` library

6704 **Table 13-34 librt Definition**

Library:	librt
SONAME:	<code>librt.so.1</code>

6706 The behavior of the interfaces in this library is specified by the following specifications:  
 6707  
 6708 [SUSv3] ISO POSIX (2003)

### 13.17.1 Shared Memory Objects

#### 13.17.1.1 Interfaces for Shared Memory Objects

An LSB conforming implementation shall provide the generic functions for Shared Memory Objects specified in Table 13-35, with the full mandatory functionality as described in the referenced underlying specification.

**Table 13-35 librt - Shared Memory Objects Function Interfaces**

shm_open [SUSv3]	shm_unlink [SUSv3]		
---------------------	-----------------------	--	--

### 13.17.2 Clock

#### 13.17.2.1 Interfaces for Clock

An LSB conforming implementation shall provide the generic functions for Clock specified in Table 13-36, with the full mandatory functionality as described in the referenced underlying specification.

**Table 13-36 librt - Clock Function Interfaces**

clock_getcpu clock id [SUSv3]	clock_getres [SUSv3]	clock_gettime [SUSv3]	clock_nanosleep [SUSv3]
clock_settime [SUSv3]			

### 13.17.3 Timers

#### 13.17.3.1 Interfaces for Timers

An LSB conforming implementation shall provide the generic functions for Timers specified in Table 13-37, with the full mandatory functionality as described in the referenced underlying specification.

**Table 13-37 librt - Timers Function Interfaces**

timer_create [SUSv3]	timer_delete [SUSv3]	timer_getoverrun [SUSv3]	timer_gettime [SUSv3]
timer_settime [SUSv3]			

### 13.18 Interfaces for libcrypt

Table 13-38 defines the library name and shared object name for the libcrypt library

**Table 13-38 libcrypt Definition**

Library:	libcrypt
----------	----------

6729	SONAME: libcrypt.so.1
------	--------------------------

6730 The behavior of the interfaces in this library is specified by the following specifications:  
 6731

6732 [SUSv3] ISO POSIX (2003)

### 13.18.1 Encryption

#### 13.18.1.1 Interfaces for Encryption

6734 An LSB conforming implementation shall provide the generic functions for  
 6735 Encryption specified in Table 13-39, with the full mandatory functionality as  
 6736 described in the referenced underlying specification.

6737 **Table 13-39 libcrypt - Encryption Function Interfaces**

6738 crypt [SUSv3]	encrypt [SUSv3]	setkey [SUSv3]	
--------------------	-----------------	----------------	--

### 13.19 Interfaces for libpam

6739 Table 13-40 defines the library name and shared object name for the libpam library

6740 **Table 13-40 libpam Definition**

6741 Library:	libpam
SONAME:	libpam.so.0

6742 The Pluggable Authentication Module (PAM) interfaces allow applications to  
 6743 request authentication via a system administrator defined mechanism, known as a  
 6744 *service*.

6745 A single service name, *other*, shall always be present. The behavior of this service  
 6746 shall be determined by the system administrator. Additional service names may also  
 6747 exist.

6748 **Note:** Future versions of this specification might define additional service names.

6749 The behavior of the interfaces in this library is specified by the following specifications:  
 6750

6751 [LSB] This Specification

### 13.19.1 Pluggable Authentication API

#### 13.19.1.1 Interfaces for Pluggable Authentication API

6753 An LSB conforming implementation shall provide the generic functions for  
 6754 Pluggable Authentication API specified in Table 13-41, with the full mandatory  
 6755 functionality as described in the referenced underlying specification.

6756 **Table 13-41 libpam - Pluggable Authentication API Function Interfaces**

pam_acct_mgmt [LSB]	pam_authenticate [LSB]	pam_chauthtok [LSB]	pam_close_session [LSB]
pam_end [LSB]	pam_fail_delay	pam_get_item	pam_getenvlist

	[LSB]	[LSB]	[LSB]
6757	pam_open_session [LSB]	pam_set_item [LSB]	pam_setcred [LSB]
	pam_strerror [LSB]		

## 13.20 Data Definitions for libpam

This section defines global identifiers and their values that are associated with interfaces contained in libpam. These definitions are organized into groups that correspond to system headers. This convention is used as a convenience for the reader, and does not imply the existence of these headers, or their content. Where an interface is defined as requiring a particular system header file all of the data definitions for that system header file presented here shall be in effect.

This section gives data definitions to promote binary application portability, not to repeat source interface definitions available elsewhere. System providers and application developers should use this ABI to supplement - not to replace - source interface definition specifications.

This specification uses the ISO C (1999) C Language as the reference programming language, and data definitions are specified in ISO C format. The C language is used here as a convenient notation. Using a C language description of these data objects does not preclude their use by other programming languages.

### 13.20.1 security/pam\_appl.h

```

6772
6773     typedef struct pam_handle pam_handle_t;
6774     struct pam_message {
6775         int msg_style;
6776         const char *msg;
6777     };
6778     struct pam_response {
6779         char *resp;
6780         int resp_retcode;
6781     };
6782
6783     struct pam_conv {
6784         int (*conv) (int num_msg, const struct pam_message **msg,
6785                      struct pam_response **resp, void *appdata_ptr);
6786         void *appdata_ptr;
6787     };
6788
6789     #define PAM_PROMPT_ECHO_OFF      1
6790     #define PAM_PROMPT_ECHO_ON       2
6791     #define PAM_ERROR_MSG           3
6792     #define PAM_TEXT_INFO            4
6793
6794     #define PAM_SERVICE              1
6795     #define PAM_USER                 2
6796     #define PAM_TTY                  3
6797     #define PAM_RHOST                4
6798     #define PAM_CONV                 5
6799     #define PAM_RUSER                8
6800     #define PAM_USER_PROMPT          9
6801
6802     #define PAM_SUCCESS               0
6803     #define PAM_OPEN_ERR              1

```

```

6804 #define PAM_USER_UNKNOWN          10
6805 #define PAM_MAXTRIES            11
6806 #define PAM_NEW_AUTHTOK_REQD    12
6807 #define PAM_ACCT_EXPIRED        13
6808 #define PAM_SESSION_ERR         14
6809 #define PAM_CRED_UNAVAIL       15
6810 #define PAM_CRED_EXPIRED        16
6811 #define PAM_CRED_ERR            17
6812 #define PAM_CONV_ERR             19
6813 #define PAM_SYMBOL_ERR           2
6814 #define PAM_AUTHTOK_ERR          20
6815 #define PAM_AUTHTOK_RECOVER_ERR  21
6816 #define PAM_AUTHTOK_LOCK_BUSY    22
6817 #define PAM_AUTHTOK_DISABLE_AGING 23
6818 #define PAM_TRY AGAIN            24
6819 #define PAM_ABORT                26
6820 #define PAM_AUTHTOK_EXPIRED      27
6821 #define PAM_BAD_ITEM              29
6822 #define PAM_SERVICE_ERR           3
6823 #define PAM_SYSTEM_ERR             4
6824 #define PAM_BUF_ERR                5
6825 #define PAM_PERM_DENIED            6
6826 #define PAM_AUTH_ERR                 7
6827 #define PAM_CRED_INSUFFICIENT     8
6828 #define PAM_AUTHINFO_UNAVAIL       9
6829
6830 #define PAM_DISALLOW_NULL_AUTHTOK 0x0001U
6831 #define PAM_ESTABLISH_CRED        0x0002U
6832 #define PAM_DELETE_CRED            0x0004U
6833 #define PAM_REINITIALIZE_CRED     0x0008U
6834 #define PAM_REFRESH_CRED           0x0010U
6835 #define PAM_CHANGE_EXPIRED_AUTHTOK 0x0020U
6836 #define PAM_SILENT                  0x8000U
6837
6838 extern int pam_set_item(pam_handle_t *, int, const void *);
6839 extern int pam_get_item(const pam_handle_t *, int, const void **);
6840 extern const char *pam_strerror(pam_handle_t *, int);
6841 extern char **pam_getenvlist(pam_handle_t *);
6842 extern int pam_fail_delay(pam_handle_t *, unsigned int);
6843 extern int pam_start(const char *, const char *, const struct pam_conv *,
6844                      pam_handle_t * * );
6845 extern int pam_end(pam_handle_t *, int);
6846 extern int pam_authenticate(pam_handle_t *, int);
6847 extern int pam_setcred(pam_handle_t *, int);
6848 extern int pam_acct_mgmt(pam_handle_t *, int);
6849 extern int pam_open_session(pam_handle_t *, int);
6850 extern int pam_close_session(pam_handle_t *, int);
6851 extern int pam_chauthtok(pam_handle_t *, int);
6852

```

## 13.21 Interface Definitions for libpam

6853 The interfaces defined on the following pages are included in libpam and are  
 6854 defined by this specification. Unless otherwise noted, these interfaces shall be  
 6855 included in the source standard.

6856 Other interfaces listed in Section 13.19 shall behave as described in the referenced  
 6857 base document.

## pam\_acct\_mgmt

### Name

6858           pam\_acct\_mgmt — establish the status of a user's account

### Synopsis

```
6859 #include <security/pam_appl.h>
6860 int pam_acct_mgmt(pam_handle_t * pamh, int flags);
```

### Description

6861           *pam\_acct\_mgmt()* establishes the account's usability and the user's accessibility to  
6862           the system. It is typically called after the user has been authenticated.

6863           *flags* may be specified as any valid flag (namely, one of those applicable to the  
6864           *flags* argument of *pam\_authenticate()*). Additionally, the value of *flags* may be  
6865           logically or'd with *PAM\_SILENT*.

### Return Value

6866           PAM\_SUCCESS

6867           Success.

6868           PAM\_NEW\_AUTHTOK\_REQD

6869           User is valid, but user's authentication token has expired. The correct response  
6870           to this return-value is to require that the user satisfy the *pam\_chauthtok()*  
6871           function before obtaining service. It may not be possible for an application to do  
6872           this. In such a case, the user should be denied access until the account password  
6873           is updated.

6874           PAM\_ACCT\_EXPIRED

6875           User is no longer permitted access to the system.

6876           PAM\_AUTH\_ERR

6877           Authentication error.

6878           PAM\_PERM\_DENIED

6879           User is not permitted to gain access at this time.

6880           PAM\_USER\_UNKNOWN

6881           User is not known to a module's account management component.

6882           **Note:** Errors may be translated to text with *pam\_strerror()*.

## pam\_authenticate

### Name

6883 pam\_authenticate — authenticate the user

### Synopsis

```
6884 #include <security/pam_appl.h>
6885 int pam_authenticate(pam_handle_t * pamh, int flags);
```

### Description

6886 pam\_authenticate() serves as an interface to the authentication mechanisms of the  
6887 loaded modules.

6888 *flags* is an optional parameter that may be specified by the following value:

6889 PAM\_DISALLOW\_NULL\_AUTHTOK

6890 Instruct the authentication modules to return PAM\_AUTH\_ERR if the user does not  
6891 have a registered authorization token.

6892 Additionally, the value of *flags* may be logically or'd with PAM\_SILENT.

6893 The process may need to be privileged in order to successfully call this function.

### Return Value

6894 PAM\_SUCCESS

6895 Success.

6896 PAM\_AUTH\_ERR

6897 User was not authenticated or process did not have sufficient privileges to  
6898 perform authentication.

6899 PAM\_CRED\_INSUFFICIENT

6900 Application does not have sufficient credentials to authenticate the user.

6901 PAM\_AUTHINFO\_UNAVAIL

6902 Modules were not able to access the authentication information. This might be  
6903 due to a network or hardware failure, etc.

6904 PAM\_USER\_UNKNOWN

6905 Supplied username is not known to the authentication service.

6906 PAM\_MAXTRIES

6907 One or more authentication modules has reached its limit of tries authenticating  
6908 the user. Do not try again.

6909 PAM\_ABORT

6910 One or more authentication modules failed to load.

6911 **Note:** Errors may be translated to text with pam\_strerror().

## pam\_chauthtok

### Name

6912 pam\_chauthtok — change the authentication token for a given user

### Synopsis

```
6913 #include <security/pam_appl.h>
6914 int pam_chauthtok(pam_handle_t * pamh, const int flags);
```

### Description

6915 pam\_chauthtok() is used to change the authentication token for a given user as  
6916 indicated by the state associated with the handle *pamh*.

6917 *flags* is an optional parameter that may be specified by the following value:

6918 PAM\_CHANGE\_EXPIRED\_AUTHTOK

6919 User's authentication token should only be changed if it has expired.

6920 Additionally, the value of *flags* may be logically or'd with PAM\_SILENT.

### RETURN VALUE

6921 PAM\_SUCCESS

6922 Success.

6923 PAM\_AUTHTOK\_ERR

6924 A module was unable to obtain the new authentication token.

6925 PAM\_AUTHTOK\_RECOVER\_ERR

6926 A module was unable to obtain the old authentication token.

6927 PAM\_AUTHTOK\_LOCK\_BUSY

6928 One or more modules were unable to change the authentication token since it is  
6929 currently locked.

6930 PAM\_AUTHTOK\_DISABLE\_AGING

6931 Authentication token aging has been disabled for at least one of the modules.

6932 PAM\_PERM\_DENIED

6933 Permission denied.

6934 PAM\_TRY AGAIN

6935 Not all modules were in a position to update the authentication token(s). In  
6936 such a case, none of the user's authentication tokens are updated.

6937 PAM\_USER\_UNKNOWN

6938 User is not known to the authentication token changing service.

6939 **Note:** Errors may be translated to text with pam\_strerror().

## pam\_close\_session

### Name

6940 pam\_close\_session — indicate that an authenticated session has ended

### Synopsis

```
6941 #include <security/pam_appl.h>
6942 int pam_close_session(pam_handle_t * pamh, int flags);
```

### Description

6943 pam\_close\_session() is used to indicate that an authenticated session has ended. It  
 6944 is used to inform the module that the user is exiting a session. It should be possible  
 6945 for the PAM library to open a session and close the same session from different  
 6946 applications.

6947 *flags* may have the value PAM\_SILENT to indicate that no output should be  
 6948 generated as a result of this function call.

### Return Value

6949 PAM\_SUCCESS

6950       Success.

6951 PAM\_SESSION\_ERR

6952       One of the required loaded modules was unable to close a session for the user.

6953       **Note:** Errors may be translated to text with pam\_strerror().

## pam\_end

### Name

6954 pam\_end — terminate the use of the PAM library

### Synopsis

```
6955 #include <security/pam_appl.h>
6956 int pam_end(pam_handle_t * pamh, int pam_status);
```

### Description

6957 pam\_end() terminates use of the PAM library. On success, the contents of \**pamh* are  
 6958 no longer valid, and all memory associated with it is invalid.

6959 Normally, *pam\_status* is passed the value PAM\_SUCCESS, but in the event of an  
 6960 unsuccessful service application, the appropriate PAM error return value should be  
 6961 used.

### Return Value

6962 PAM\_SUCCESS

6963       Success.

6964       **Note:** Errors may be translated to text with pam\_strerror().

## pam\_fail\_delay

### Name

6965 pam\_fail\_delay — specify delay time to use on authentication error

### Synopsis

6966 #include <security/pam\_appl.h>  
6967 int pam\_fail\_delay(pam\_handle\_t \* *pamh*, unsigned int *micro\_sec*);

### Description

6968 pam\_fail\_delay() specifies the minimum delay for the PAM library to use when  
6969 an authentication error occurs. The actual delay can vary by as much at 25%. If this  
6970 function is called multiple times, the longest time specified by any of the call will be  
6971 used.

6972 The delay is invoked if an authentication error occurs during the  
6973 pam\_authenticate() or pam\_chauthtok() function calls.

6974 Independent of the success of pam\_authenticate() or pam\_chauthtok(), the delay  
6975 time is reset to its default value of 0 when the PAM library returns control to the  
6976 application from these two functions.

### Return Value

6977 PAM\_SUCCESS

6978 Success.

6979 **Note:** Errors may be translated to text with pam\_strerror().

## pam\_get\_item

### Name

6980 pam\_get\_item — obtain the value of the indicated item.

### Synopsis

```
6981 #include <security/pam_appl.h>
6982 int pam_get_item(const pam_handle_t * pamh, int item_type, const void * *
6983 item);
```

### Description

6984 pam\_get\_item() obtains the value of the indicated *item\_type*. The possible values  
 6985 of *item\_type* are the same as listed for pam\_set\_item().

6986 On success, *item* contains a pointer to the value of the corresponding item. Note that  
 6987 this is a pointer to the actual data and should not be free()'d or over-written.

### Return Value

6988 PAM\_SUCCESS

6989 Success.

6990 PAM\_PERM\_DENIED

6991 Application passed a NULL pointer for *item*.

6992 PAM\_BAD\_ITEM

6993 Application attempted to get an undefined item.

6994 **Note:** Errors may be translated to text with pam\_strerror().

## pam\_getenvlist

### Name

6995 pam\_getenvlist — returns a pointer to the complete PAM environment.

### Synopsis

```
6996 #include <security/pam_appl.h>
6997 char * const * pam_getenvlist(pam_handle_t * pamh);
```

### Description

6998 pam\_getenvlist() returns a pointer to the complete PAM environment. This  
 6999 pointer points to an array of pointers to NUL-terminated strings and must be  
 7000 terminated by a NULL pointer. Each string has the form "name=value".

7001 The PAM library module allocates memory for the returned value and the  
 7002 associated strings. The calling application is responsible for freeing this memory.

### Return Value

7003 pam\_getenvlist() returns an array of string pointers containing the PAM  
 7004 environment. On error, NULL is returned.

## pam\_open\_session

### Name

7005 pam\_open\_session — indicate session has started

### Synopsis

7006 #include <security/pam\_appl.h>  
7007 int pam\_open\_session(pam\_handle\_t \* pamh, int flags);

### Description

7008 The `pam_open_session()` function is used to indicate that an authenticated session  
7009 has begun, after the user has been identified (see `pam_authenticate()`) and, if  
7010 necessary, granted credentials (see `pam_setcred()`). It is used to inform the module  
7011 that the user is currently in a session. It should be possible for the PAM library to  
7012 open a session and close the same session from different applications.

7013 `flags` may have the value `PAM_SILENT` to indicate that no output be generated as a  
7014 result of this function call.

### Return Value

7015 PAM\_SUCCESS

7016 Success.

7017 PAM\_SESSION\_ERR

7018 One of the loaded modules was unable to open a session for the user.

7019 **Note:** Errors may be translated to text with `pam_strerror()`.

## pam\_set\_item

### Name

7020 pam\_set\_item — (re)set the value of an item.

### Synopsis

```
7021 #include <security/pam_appl.h>
7022 int pam_set_item(pam_handle_t * pamh, int item_type, const void * item);
```

### Description

7023 pam\_set\_item( ) (re)sets the value of one of the following item\_types:

7024 PAM\_SERVICE  
7025 service name

7026 PAM\_USER  
7027 user name

7028 PAM\_TTY  
7029 terminal name

7030 The value for a device file should include the /dev/ prefix. The value for  
7031 graphical, X-based, applications should be the \$DISPLAY variable.

7032 PAM\_RHOST  
7033 remote host name

7034 PAM\_CONV  
7035 conversation structure

7036 PAM\_RUSER  
7037 remote user name

7038 PAM\_USER\_PROMPT  
7039 string to be used when prompting for a user's name  
7040 The default value for this string is Please enter username: .

7041 For all *item\_types* other than PAM\_CONV, *item* is a pointer to a NULL-terminated  
7042 character string. In the case of PAM\_CONV, *item* points to an initialized pam\_conv  
7043 structure.

### Return Value

7044 PAM\_SUCCESS  
7045 Success.

7046 PAM\_PERM\_DENIED  
7047 An attempt was made to replace the conversation structure with a NULL value.

7048 PAM\_BUF\_ERR

7049                   Function ran out of memory making a copy of the item.

7050                   PAM\_BAD\_ITEM

7051                   Application attempted to set an undefined item.

7052                   **Note:** Errors may be translated to text with `pam_strerror()`.

## pam\_setcred

### Name

7053 pam\_setcred — set the module-specific credentials of the user

### Synopsis

```
7054 #include <security/pam_appl.h>
7055 extern int pam_setcred(pam_handle_t * pamh, int flags);
```

### Description

7056 pam\_setcred( ) sets the module-specific credentials of the user. It is usually called  
 7057 after the user has been authenticated, after the account management function has  
 7058 been called and after a session has been opened for the user.

7059 *flags* maybe specified from among the following values:

7060 PAM\_ESTABLISH\_CRED  
 7061       set credentials for the authentication service

7062 PAM\_DELETE\_CRED  
 7063       delete credentials associated with the authentication service

7064 PAM\_REINITIALIZE\_CRED  
 7065       reinitialize the user credentials

7066 PAM\_REFRESH\_CRED  
 7067       extend lifetime of the user credentials

7068 Additionally, the value of *flags* may be logically or'd with PAM\_SILENT.

### Return Value

7069 PAM\_SUCCESS  
 7070       Success.

7071 PAM\_CRED\_UNAVAIL  
 7072       Module cannot retrieve the user's credentials.

7073 PAM\_CRED\_EXPIRED  
 7074       User's credentials have expired.

7075 PAM\_USER\_UNKNOWN  
 7076       User is not known to an authentication module.

7077 PAM\_CRED\_ERR  
 7078       Module was unable to set the credentials of the user.

7079 **Note:** Errors may be translated to text with pam\_strerror( ).

## pam\_start

### Name

7080 pam\_start — initialize the PAM library

### Synopsis

```
7081 #include <security/pam_appl.h>
7082 int pam_start(const char * service_name, const char * user, const struct
7083 pam_conv * pam_conversation, pam_handle_t ** pamh);
```

### Description

7084 pam\_start() is used to initialize the PAM library. It must be called prior to any  
 7085 other usage of the PAM library. On success, \**pamh* becomes a handle that provides  
 7086 continuity for successive calls to the PAM library. pam\_start() expects arguments  
 7087 as follows: the *service\_name* of the program, the *username* of the individual to be  
 7088 authenticated, a pointer to an application-supplied *pam\_conv* structure, and a  
 7089 pointer to a *pam\_handle\_t* pointer.

7090 An application must provide the *conversation function* used for direct communication  
 7091 between a loaded module and the application. The application also typically  
 7092 provides a means for the module to prompt the user for a password, etc.

7093 The structure, *pam\_conv*, is defined to be,

```
7094 struct pam_conv {
7095     int (*conv) (int num_msg,
7096                  const struct pam_message * *msg,
7097                  struct pam_response * *resp,
7098                  void *appdata_ptr);
7099     void *appdata_ptr;
```

7100 } ;

7101 It is initialized by the application before it is passed to the library. The contents of  
 7102 this structure are attached to the *\*pamh* handle. The point of this argument is to  
 7103 provide a mechanism for any loaded module to interact directly with the application  
 7104 program; this is why it is called a conversation structure.

7105 When a module calls the referenced `conv()` function, *appdata\_ptr* is set to the  
 7106 second element of this structure.

7107 The other arguments of a call to `conv()` concern the information exchanged by  
 7108 module and application. *num\_msg* holds the length of the array of pointers passed via  
 7109 *msg*. On success, the pointer *resp* points to an array of *num\_msg pam\_response*  
 7110 structures, holding the application-supplied text. Note that *resp* is a struct  
 7111 *pam\_response* array and not an array of pointers.

## Return Value

7112 PAM\_SUCCESS

7113 Success.

7114 PAM\_BUF\_ERR

7115 Memory allocation error.

7116 PAM\_ABORT

7117 Internal failure.

## ERRORS

7118 May be translated to text with `pam_strerror()`.

## **pam\_strerror**

### Name

7119 `pam_strerror` — returns a string describing the PAM error

### Synopsis

```
7120 #include <security/pam_appl.h>
7121 const char * pam_strerror(pam_handle_t * pamh, int errnum);
```

### Description

7122 `pam_strerror()` returns a string describing the PAM error associated with *errnum*.

### Return Value

7123 On success, this function returns a description of the indicated error. The application  
 7124 should not free or modify this string. Otherwise, a string indicating that the error is  
 7125 unknown shall be returned. It is unspecified whether or not the string returned is  
 7126 translated according to the setting of LC\_MESSAGES.

## **IV Utility Libraries**

## 14 Utility Libraries

### 14.1 Introduction

1 An LSB-conforming implementation shall also support the following utility libraries  
2 which are built on top of the interfaces provided by the base libraries. These libraries  
3 implement common functionality, and hide additional system dependent  
4 information such as file formats and device names.

- 5 • libz  
6 • libcurses  
7 • libutil

8 The structure of the definitions for these libraries follows the same model as used for  
9 Base Libraries.

### 14.2 Interfaces for libz

10 Table 14-1 defines the library name and shared object name for the libz library

11 **Table 14-1 libz Definition**

Library:	libz
SONAME:	libz.so.1

13 The behavior of the interfaces in this library is specified by the following specifications:  
14

15 [LSB] This Specification

#### 14.2.1 Compression Library

##### 14.2.1.1 Interfaces for Compression Library

17 An LSB conforming implementation shall provide the generic functions for  
18 Compression Library specified in Table 14-2, with the full mandatory functionality  
19 as described in the referenced underlying specification.

20 **Table 14-2 libz - Compression Library Function Interfaces**

adler32 [LSB]	compress [LSB]	compress2 [LSB]	compressBound [LSB]
crc32 [LSB]	deflate [LSB]	deflateBound [LSB]	deflateCopy [LSB]
deflateEnd [LSB]	deflateInit2_ [LSB]	deflateInit_ [LSB]	deflateParams [LSB]
deflateReset [LSB]	deflateSetDictionary [LSB]	get_crc_table [LSB]	gzclose [LSB]
gzdopen [LSB]	gzeof [LSB]	gzerror [LSB]	gzflush [LSB]
gzgetc [LSB]	gzgets [LSB]	gzopen [LSB]	gzprintf [LSB]
gzputc [LSB]	gzputs [LSB]	gzread [LSB]	gzrewind [LSB]

21

gzseek [LSB]	gzsetparams [LSB]	gztell [LSB]	gzwrite [LSB]
inflate [LSB]	inflateEnd [LSB]	inflateInit2_ [LSB]	inflateInit_ [LSB]
inflateReset [LSB]	inflateSetDictionary [LSB]	inflateSync [LSB]	inflateSyncPoint [LSB]
uncompress [LSB]	zError [LSB]	zlibVersion [LSB]	

## 14.3 Data Definitions for libz

This section defines global identifiers and their values that are associated with interfaces contained in libz. These definitions are organized into groups that correspond to system headers. This convention is used as a convenience for the reader, and does not imply the existence of these headers, or their content. Where an interface is defined as requiring a particular system header file all of the data definitions for that system header file presented here shall be in effect.

This section gives data definitions to promote binary application portability, not to repeat source interface definitions available elsewhere. System providers and application developers should use this ABI to supplement - not to replace - source interface definition specifications.

This specification uses the ISO C (1999) C Language as the reference programming language, and data definitions are specified in ISO C format. The C language is used here as a convenient notation. Using a C language description of these data objects does not preclude their use by other programming languages.

### 14.3.1 zlib.h

In addition to the values below, the `zlib.h` header shall define the `ZLIB_VERSION` macro. This macro may be used to check that the version of the library at run time matches that at compile time.

See also the `zlibVersion()` function, which returns the library version at run time. The first character of the version at compile time should always match the first character at run time.

```

42 #define Z_NULL    0
43 #define MAX_WBITS      15
44 #define MAX_MEM_LEVEL   9
45 #define deflateInit2(strm,level,method>windowBits,memLevel,strategy)
46 \
47
48
49     deflateInit2((strm),(level),(method),(windowBits),(memLevel),(strat
50     egy),ZLIB_VERSION,sizeof(z_stream))
51 #define deflateInit(strm,level) \
52     deflateInit_((strm), (level),           ZLIB_VERSION,
53     sizeof(z_stream))
54 #define inflateInit2(strm>windowBits) \
55     inflateInit2_((strm), (windowBits), ZLIB_VERSION,
56     sizeof(z_stream))
57 #define inflateInit(strm) \
58     inflateInit_((strm),           ZLIB_VERSION, sizeof(z_stream))
59
60     typedef char charf;
61     typedef int intf;
62

```

```

63     typedef void *voidpf;
64     typedef unsigned int uInt;
65     typedef unsigned long int uLong;
66     typedef uLong uLongf;
67     typedef void *voidp;
68     typedef unsigned char Byte;
69     typedef off_t z_off_t;
70     typedef void *const voidpc;
71
72     typedef voidpf(*alloc_func) (voidpf opaque, uInt items, uInt size);
73     typedef void (*free_func) (voidpf opaque, voidpf address);
74     struct internal_state {
75         int dummy;
76     };
77     typedef Byte Bytef;
78     typedef uInt uIntf;
79
80     typedef struct z_stream_s {
81         Bytef *next_in;
82         uInt avail_in;
83         uLong total_in;
84         Bytef *next_out;
85         uInt avail_out;
86         uLong total_out;
87         char *msg;
88         struct internal_state *state;
89         alloc_func zalloc;
90         free_func zfree;
91         voidpf opaque;
92         int data_type;
93         uLong adler;
94         uLong reserved;
95     } z_stream;
96
97     typedef z_stream *z_streamp;
98     typedef voidp gzFile;
99
100    #define Z_NO_FLUSH      0
101    #define Z_PARTIAL_FLUSH 1
102    #define Z_SYNC_FLUSH    2
103    #define Z_FULL_FLUSH    3
104    #define Z_FINISH        4
105
106    #define Z_ERRNO          (-1)
107    #define Z_STREAM_ERROR   (-2)
108    #define Z_DATA_ERROR     (-3)
109    #define Z_MEM_ERROR      (-4)
110    #define Z_BUF_ERROR      (-5)
111    #define Z_VERSION_ERROR (-6)
112    #define Z_OK             0
113    #define Z_STREAM_END     1
114    #define Z_NEED_DICT      2
115
116    #define Z_DEFAULT_COMPRESSION (-1)
117    #define Z_NO_COMPRESSION  0
118    #define Z_BEST_SPEED     1
119    #define Z_BEST_COMPRESSION 9
120
121    #define Z_DEFAULT_STRATEGY 0
122    #define Z_FILTERED        1
123    #define Z_HUFFMAN_ONLY   2
124
125    #define Z_BINARY          0
126    #define Z_ASCII           1

```

```

127 #define Z_UNKNOWN          2
128 #define Z_DEFLATED         8
129
130 extern int gzread(gzFile, voidp, unsigned int);
131 extern int gzclose(gzFile);
132 extern gzFile gzopen(const char *, const char *);
133 extern gzFile gdopen(int, const char *);
134 extern int gzwrite(gzFile, voidpc, unsigned int);
135 extern int gflush(gzFile, int);
136 extern const char *gzerror(gzFile, int *);
137 extern uLong adler32(uLong, const Bytef *, UInt);
138 extern int compress(Bytef *, uLongf *, const Bytef *, uLong);
139 extern int compress2(Bytef *, uLongf *, const Bytef *, uLong, int);
140 extern uLong crc32(uLong, const Bytef *, UInt);
141 extern int deflate(z_streamp, int);
142 extern int deflateCopy(z_streamp, z_streamp);
143 extern int deflateEnd(z_streamp);
144 extern int deflateInit2_(z_streamp, int, int, int, int, int, const char
145   *,
146   int);
147 extern int deflateInit_(z_streamp, int, const char *, int);
148 extern int deflateParams(z_streamp, int, int);
149 extern int deflateReset(z_streamp);
150 extern int deflateSetDictionary(z_streamp, const Bytef *, UInt);
151 extern const uLongf *get_crc_table(void);
152 extern int gzeof(gzFile);
153 extern int gzgetc(gzFile);
154 extern char *gzgets(gzFile, char *, int);
155 extern int gzprintf(gzFile, const char *, ...);
156 extern int gzputc(gzFile, int);
157 extern int gzputs(gzFile, const char *);
158 extern int gzrewind(gzFile);
159 extern z_off_t gzseek(gzFile, z_off_t, int);
160 extern int gzsetparams(gzFile, int, int);
161 extern z_off_t gtell(gzFile);
162 extern int inflate(z_streamp, int);
163 extern int inflateEnd(z_streamp);
164 extern int inflateInit2_(z_streamp, int, const char *, int);
165 extern int inflateInit_(z_streamp, const char *, int);
166 extern int inflateReset(z_streamp);
167 extern int inflateSetDictionary(z_streamp, const Bytef *, UInt);
168 extern int inflateSync(z_streamp);
169 extern int inflateSyncPoint(z_streamp);
170 extern int uncompress(Bytef *, uLongf *, const Bytef *, uLong);
171 extern const char *zError(int);
172 extern const char *zlibVersion(void);
173 extern uLong deflateBound(z_streamp, uLong);
174 extern uLong compressBound(uLong);
175

```

## 14.4 Interface Definitions for libz

176 The interfaces defined on the following pages are included in libz and are defined by  
 177 this specification. Unless otherwise noted, these interfaces shall be included in the  
 178 source standard.

179 Other interfaces listed in Section 14.2 shall behave as described in the referenced  
 180 base document.

## adler32

### Name

181 adler32 — compute Adler 32 Checksum

### Synopsis

```
182 #include <zlib.h>
183 uLong adler32(uLong adler, const Bytef * buf, uInt len);
```

### Description

184 The `adler32()` function shall compute a running Adler-32 checksum (as described  
 185 in RFC 1950: ZLIB Compressed Data Format Specification). On entry, `adler` is the  
 186 previous value for the checksum, and `buf` shall point to an array of `len` bytes of data  
 187 to be added to this checksum. The `adler32()` function shall return the new  
 188 checksum.

189 If `buf` is `NULL` (or `Z_NULL`), `adler32()` shall return the initial checksum.

### Return Value

190 The `adler32()` function shall return the new checksum value.

### Errors

191 None defined.

### Application Usage (informative)

192 The following code fragment demonstrates typical usage of the `adler32()` function:

```
193     uLong adler = adler32(0L, Z_NULL, 0);
194
195     while (read_buffer(buffer, length) != EOF) {
196         adler = adler32(adler, buffer, length);
197     }
198     if (adler != original_adler) error();
```

## compress

### Name

199 compress — compress data

### Synopsis

```
200 #include <zlib.h>
201 int compress(Bytef * dest, uLongf * destLen, const Bytef * source, uLong
202 sourceLen);
```

### Description

203 The `compress()` function shall attempt to compress `sourceLen` bytes of data in the  
204 buffer `source`, placing the result in the buffer `dest`.

205 On entry, `destLen` should point to a value describing the size of the `dest` buffer. The  
206 application should ensure that this value be at least  $(\text{sourceLen} \times 1.001) + 12$ . On  
207 successful exit, the variable referenced by `destLen` shall be updated to hold the  
208 length of compressed data in `dest`.

209 The `compress()` function is equivalent to `compress2()` with a `level` of  
210 `Z_DEFAULT_LEVEL`.

### Return Value

211 On success, `compress()` shall return `Z_OK`. Otherwise, `compress()` shall return a  
212 value to indicate the error.

### Errors

213 On error, `compress()` shall return a value as described below:

214 `Z_BUF_ERROR`

215     The buffer `dest` was not large enough to hold the compressed data.

216 `Z_MEM_ERROR`

217     Insufficient memory.

## compress2

### Name

218 compress2 — compress data at a specified level

### Synopsis

```
219 #include <zlib.h>
220 int compress2(Bytef * dest, uLongf * destLen, const Bytef * source, uLong
221 sourceLen, int level);
```

### Description

222 The `compress2()` function shall attempt to compress `sourceLen` bytes of data in the  
 223 buffer `source`, placing the result in the buffer `dest`, at the level described by `level`.  
 224 The `level` supplied shall be a value between 0 and 9, or the value  
 225 `Z_DEFAULT_COMPRESSION`. A `level` of 1 requests the highest speed, while a `level` of  
 226 9 requests the highest compression. A `level` of 0 indicates that no compression  
 227 should be used, and the output shall be the same as the input.

228 On entry, `destLen` should point to a value describing the size of the `dest` buffer. The  
 229 application should ensure that this value be at least  $(sourceLen \times 1.001) + 12$ . On  
 230 successful exit, the variable referenced by `destLen` shall be updated to hold the  
 231 length of compressed data in `dest`.

232 The `compress()` function is equivalent to `compress2()` with a `level` of  
 233 `Z_DEFAULT_LEVEL`.

### Return Value

234 On success, `compress2()` shall return `Z_OK`. Otherwise, `compress2()` shall return a  
 235 value to indicate the error.

### Errors

236 On error, `compress2()` shall return a value as described below:

237 `Z_BUF_ERROR`

238     The buffer `dest` was not large enough to hold the compressed data.

239 `Z_MEM_ERROR`

240     Insufficient memory.

241 `Z_STREAM_ERROR`

242     The `level` was not `Z_DEFAULT_LEVEL`, or was not between 0 and 9.

## compressBound

### Name

243      compressBound — compute compressed data size

### Synopsis

244      #include <zlib.h>  
245      int compressBound(uLong sourceLen);

### Description

246      The `compressBound()` function shall estimate the size of buffer required to  
247      compress `sourceLen` bytes of data using the `compress()` or `compress2()` functions.  
248      If successful, the value returned shall be an upper bound for the size of buffer  
249      required to compress `sourceLen` bytes of data, using the parameters stored in  
250      *stream*, in a single call to `compress()` or `compress2()`.

### Return Value

251      The `compressBound()` shall return a value representing the upper bound of an array  
252      to allocate to hold the compressed data in a single call to `compress()` or  
253      `compress2()`. This function may return a conservative value that may be larger than  
254      `sourceLen`.

### Errors

255      None defined.

## crc32

### Name

256           crc32 — compute CRC-32 Checksum

### Synopsis

```
257 #include <zlib.h>
258 uLong crc32(uLong crc, const Bytef * buf, uint len);
```

### Description

259           The `crc32()` function shall compute a running Cyclic Redundancy Check checksum,  
 260           as defined in ITU-T V.42. On entry, `crc` is the previous value for the checksum, and  
 261           `buf` shall point to an array of `len` bytes of data to be added to this checksum. The  
 262           `crc32()` function shall return the new checksum.

263           If `buf` is `NULL` (or `Z_NULL`), `crc32()` shall return the initial checksum.

### Return Value

264           The `crc32()` function shall return the new checksum value.

### Errors

265           None defined.

### Application Usage (informative)

266           The following code fragment demonstrates typical usage of the `crc32()` function:

```
267           uLong crc = crc32(0L, Z_NULL, 0);
268
269           while (read_buffer(buffer, length) != EOF) {
270                crc = crc32(crc, buffer, length);
271               }
272           if (crc != original_crc) error();
```

## deflate

### Name

273      `deflate — compress data`

### Synopsis

274      `#include <zlib.h>`  
 275      `int deflate(z_streamp stream, int flush);`

### Description

276      The `deflate()` function shall attempt to compress data until either the input buffer  
 277      is empty or the output buffer is full. The `stream` references a `z_stream` structure.  
 278      Before the first call to `deflate()`, this structure should have been initialized by a call  
 279      to `deflateInit2_()`.

280      **Note:** `deflateInit2_()` is only in the binary standard; source level applications should  
 281      initialize `stream` via a call to `deflateInit()` or `deflateInit2()`.

282      In addition, the `stream` input and output buffers should have been initialized as  
 283      follows:

284      `next_in`  
 285          should point to the data to be compressed.  
 286      `avail_in`  
 287          should contain the number of bytes of data in the buffer referenced by `next_in`.  
 288      `next_out`  
 289          should point to a buffer where compressed data may be placed.  
 290      `avail_out`  
 291          should contain the size in bytes of the buffer referenced by `next_out`

292      The `deflate()` function shall perform one or both of the following actions:

- 293      1. Compress input data from `next_in` and update `next_in`, `avail_in` and  
        `total_in` to reflect the data that has been compressed.
- 295      2. Fill the output buffer referenced by `next_out`, and update `next_out`,  
        `avail_out` and `total_out` to reflect the compressed data that has been placed  
        there. If `flush` is not `Z_NO_FLUSH`, and `avail_out` indicates that there is still  
        space in output buffer, this action shall always occur (see below for further  
        details).

300      The `deflate()` function shall return when either `avail_in` reaches zero (indicating  
 301      that all the input data has been compressed), or `avail_out` reaches zero (indicating  
 302      that the output buffer is full).

303      On success, the `deflate()` function shall set the `adler` field of the `stream` to the  
 304      `adler32()` checksum of all the input data compressed so far (represented by  
 305      `total_in`).

306 If the `deflate()` function shall attempt to determine the type of input data, and set  
 307 field `data_type` in `stream` to `z_ASCII` if the majority of the data bytes fall within the  
 308 ASCII (ISO 646) printable character range. Otherwise, it shall set `data_type` to  
 309 `z_BINARY`. This data type is informational only, and does not affect the compression  
 310 algorithm.

311 **Note:** Future versions of the LSB may remove this requirement, since it is based on an  
 312 outdated character set that does not support Internationalization, and does not affect the  
 313 algorithm. It is included for information only at this release. Applications should not  
 314 depend on this field.

## 315 **Flush Operation**

316 The parameter `flush` determines when compressed bits are added to the output  
 317 buffer in `next_out`. If `flush` is `z_NO_FLUSH`, `deflate()` may return with some data  
 318 pending output, and not yet added to the output buffer.

319 If `flush` is `z_SYNC_FLUSH`, `deflate()` shall flush all pending output to `next_out`  
 320 and align the output to a byte boundary. A synchronization point is generated in the  
 321 output.

322 If `flush` is `z_FULL_FLUSH`, all output shall be flushed, as for `z_SYNC_FLUSH`, and the  
 323 compression state shall be reset. A synchronization point is generated in the output.

324 **Rationale:** `z_SYNC_FLUSH` is intended to ensure that the compressed data contains all the  
 325 data compressed so far, and allows a decompressor to reconstruct all of the input data.  
 326 `z_FULL_FLUSH` allows decompression to restart from this point if the previous  
 327 compressed data has been lost or damaged. Flushing is likely to degrade the  
 328 performance of the compression system, and should only be used where necessary.

329 If `flush` is set to `z_FINISH`, all pending input shall be processed and `deflate()`  
 330 shall return with `Z_STREAM_END` if there is sufficient space in the output buffer at  
 331 `next_out`, as indicated by `avail_out`. If `deflate()` is called with `flush` set to  
 332 `z_FINISH` and there is insufficient space to store the compressed data, and no other  
 333 error has occurred during compression, `deflate()` shall return `Z_OK`, and the  
 334 application should call `deflate()` again with `flush` unchanged, and having  
 335 updated `next_out` and `avail_out`.

336 If all the compression is to be done in a single step, `deflate()` may be called with  
 337 `flush` set to `z_FINISH` immediately after the stream has been initialized if  
 338 `avail_out` is set to at least the value returned by `deflateBound()`.

## Return Value

339 On success, `deflate()` shall return `Z_OK`, unless `flush` was set to `z_FINISH` and  
 340 there was sufficient space in the output buffer to compress all of the input data. In  
 341 this case, `deflate()` shall return `Z_STREAM_END`. On error, `deflate()` shall  
 342 return a value to indicate the error.

343 **Note:** If `deflate()` returns `Z_OK` and has set `avail_out` to zero, the function should be  
 344 called again with the same value for `flush`, and with updated `next_out` and `avail_out`  
 345 until `deflate()` returns with `Z_OK` (or `Z_STREAM_END` if `flush` is set to `z_FINISH`)  
 346 and a non-zero `avail_out`.

## Errors

347 On error, `deflate()` shall return a value as described below, and set the `msg` field of  
 348 `stream` to point to a string describing the error:

```

349     Z_BUF_ERROR
350         No progress is possible; either avail_in or avail_out was zero.

351     Z_MEM_ERROR
352         Insufficient memory.

353     Z_STREAM_ERROR
354         The state (as represented in stream) is inconsistent, or stream was NULL.

```

## deflateBound

### Name

355 `deflateBound` — compute compressed data size

### Synopsis

```

356 #include <zlib.h>
357 int deflateBound(z_streamp stream, uLong sourceLen);

```

### Description

358 The `deflateBound()` function shall estimate the size of buffer required to compress  
 359 *sourceLen* bytes of data. If successful, the value returned shall be an upper bound  
 360 for the size of buffer required to compress *sourceLen* bytes of data, using the  
 361 parameters stored in *stream*, in a single call to `deflate()` with flush set to  
 362 `Z_FINISH`.

363 On entry, *stream* should have been initialized via a call to `deflateInit_()` or  
 364 `deflateInit2_()`.

### Return Value

365 The `deflateBound()` shall return a value representing the upper bound of an array  
 366 to allocate to hold the compressed data in a single call to `deflate()`. If the *stream* is  
 367 not correctly initialized, or is `NULL`, then `deflateBound()` may return a conservative  
 368 value that may be larger than *sourceLen*.

### Errors

369 None defined.

## deflateCopy

### Name

370           deflateCopy — copy compression stream

### Synopsis

371           #include <zlib.h>  
372           int deflateCopy(z\_streamp dest, z\_streamp source);

### Description

373           The deflateCopy() function shall copy the compression state information in  
374           *source* to the uninitialized *z\_stream* structure referenced by *dest*.

375           On successful return, *dest* will be an exact copy of the stream referenced by *source*.  
376           The input and output buffer pointers in *next\_in* and *next\_out* will reference the  
377           same data.

### Return Value

378           On success, deflateCopy() shall return Z\_OK. Otherwise it shall return a value less  
379           than zero to indicate the error.

### Errors

380           On error, deflateCopy() shall return a value as described below:

381           Z\_STREAM\_ERROR

382                 The state in *source* is inconsistent, or either *source* or *dest* was NULL.

383           Z\_MEM\_ERROR

384                 Insufficient memory available.

### Application Usage (informative)

385           This function can be useful when several compression strategies will be tried, for  
386           example when there are several ways of pre-processing the input data with a filter.  
387           The streams that will be discarded should then be freed by calling deflateEnd().  
388           Note that deflateCopy() duplicates the internal compression state which can be  
389           quite large, so this strategy may be slow and can consume lots of memory.

## deflateEnd

### Name

390      deflateEnd — free compression stream state

### Synopsis

391      #include <zlib.h>  
392      int deflateEnd(z\_streamp *stream*);

### Description

393      The deflateEnd() function shall free all allocated state information referenced by  
394      *stream*. All pending output is discarded, and unprocessed input is ignored.

### Return Value

395      On success, deflateEnd() shall return Z\_OK, or Z\_DATA\_ERROR if there was  
396      pending output discarded or input unprocessed. Otherwise it shall return  
397      Z\_STREAM\_ERROR to indicate the error.

### Errors

398      On error, deflateEnd() shall return Z\_STREAM\_ERROR. The following conditions  
399      shall be treated as an error:

- 400      • The state in *stream* is inconsistent or inappropriate.  
401      • *stream* is NULL.

## **deflateInit2\_**

### **Name**

402      `deflateInit2_ — initialize compression system`

### **Synopsis**

```
403 #include <zlib.h>
404     int deflateInit2_ (z_streamp strm, int level, int method, int windowBits,
405     int memLevel, int strategy, char * version, int stream_size);
```

### **Description**

406      The `deflateInit2_()` function shall initialize the compression system. On entry,  
 407      *strm* shall refer to a user supplied `z_stream` object (a `z_stream_s` structure). The  
 408      following fields shall be set on entry:

409      *zalloc*

410        a pointer to an `alloc_func` function, used to allocate state information. If this is  
 411        `NULL`, a default allocation function will be used.

412      *zfree*

413        a pointer to a `free_func` function, used to free memory allocated by the `zalloc`  
 414        function. If this is `NULL` a default free function will be used.

415      *opaque*

416        If `alloc_func` is not `NULL`, `opaque` is a user supplied pointer to data that will be  
 417        passed to the `alloc_func` and `free_func` functions.

418        If the `version` requested is not compatible with the version implemented, or if the  
 419        size of the `z_stream_s` structure provided in `stream_size` does not match the size  
 420        in the library implementation, `deflateInit2_()` shall fail, and return  
 421        `Z_VERSION_ERROR`.

422        The `level` supplied shall be a value between 0 and 9, or the value  
 423        `Z_DEFAULT_COMPRESSION`. A `level` of 1 requests the highest speed, while a `level` of  
 424        9 requests the highest compression. A `level` of 0 indicates that no compression  
 425        should be used, and the output shall be the same as the input.

426        The `method` selects the compression algorithm to use. LSB conforming  
 427        implementation shall support the `Z_DEFLATED` method, and may support other  
 428        implementation defined methods.

429        The `windowBits` parameter shall be a base 2 logarithm of the window size to use,  
 430        and shall be a value between 8 and 15. A smaller value will use less memory, but  
 431        will result in a poorer compression ratio, while a higher value will give better  
 432        compression but utilize more memory.

433        The `memLevel` parameter specifies how much memory to use for the internal state.  
 434        The value of `memLevel` shall be between 1 and `MAX_MEM_LEVEL`. Smaller values use  
 435        less memory but are slower, while higher values use more memory to gain  
 436        compression speed.

437        The `strategy` parameter selects the compression strategy to use:

438        `Z_DEFAULT_STRATEGY`

439               use the system default compression strategy. `Z_DEFAULT_STRATEGY` is  
 440               particularly appropriate for text data.

441     `Z_FILTERED`  
 442               use a compression strategy tuned for data consisting largely of small values  
 443               with a fairly random distribution. `Z_FILTERED` uses more Huffman encoding  
 444               and less string matching than `Z_DEFAULT_STRATEGY`.

445     `Z_HUFFMAN_ONLY`  
 446               force Huffman encoding only, with no string match.

447     The `deflateInit2_()` function is not in the source standard; it is only in the binary  
 448     standard. Source applications should use the `deflateInit2()` macro.

## Return Value

449     On success, the `deflateInit2_()` function shall return `z_OK`. Otherwise,  
 450     `deflateInit2_()` shall return a value as described below to indicate the error.

## Errors

451     On error, `deflateInit2_()` shall return one of the following error indicators:

452     `Z_STREAM_ERROR`  
 453               Invalid parameter.

454     `Z_MEM_ERROR`  
 455               Insufficient memory available.

456     `Z_VERSION_ERROR`  
 457               The version requested is not compatible with the library version, or the  
 458               `z_stream` size differs from that used by the library.

459     In addition, the `msg` field of the `strm` may be set to an error message.

## **deflateInit\_**

### **Name**

460      `deflateInit_` — initialize compression system

### **Synopsis**

```
461 #include <zlib.h>
462 int deflateInit_(z_streamp stream, int level, const char * version, int
463 stream_size);
```

### **Description**

464 The `deflateInit_()` function shall initialize the compression system. On entry,  
 465 `stream` shall refer to a user supplied `z_stream` object (a `z_stream_s` structure). The  
 466 following fields shall be set on entry:

467      `zalloc`

468      a pointer to an `alloc_func` function, used to allocate state information. If this is  
 469 `NULL`, a default allocation function will be used.

470      `zfree`

471      a pointer to a `free_func` function, used to free memory allocated by the `zalloc`  
 472 function. If this is `NULL` a default free function will be used.

473      `opaque`

474      If `alloc_func` is not `NULL`, `opaque` is a user supplied pointer to data that will be  
 475 passed to the `alloc_func` and `free_func` functions.

476      If the `version` requested is not compatible with the version implemented, or if the  
 477 size of the `z_stream_s` structure provided in `stream_size` does not match the size  
 478 in the library implementation, `deflateInit_()` shall fail, and return  
 479 `Z_VERSION_ERROR`.

480      The `level` supplied shall be a value between 0 and 9, or the value  
 481 `Z_DEFAULT_COMPRESSION`. A `level` of 1 requests the highest speed, while a `level` of  
 482 9 requests the highest compression. A `level` of 0 indicates that no compression  
 483 should be used, and the output shall be the same as the input.

484      The `deflateInit_()` function is not in the source standard; it is only in the binary  
 485 standard. Source applications should use the `deflateInit()` macro.

486      The `deflateInit_()` function is equivalent to

487      `deflateInit2_(stream, level, Z_DEFLATED, MAX_WBITS, DEF_MEM_LEVEL,`

`Z_DEFAULT_STRATEGY, version, stream_size);`

## Return Value

489           **On success**, the deflateInit\_() function shall return `z_OK`. Otherwise,  
490           deflateInit\_() shall return a value as described below to indicate the error.

## Errors

491 On error, deflateInit\_() shall return one of the following error indicators:

492 Z\_STREAM\_ERROR

Invalid parameter.

494 Z\_MEM\_ERROR

Insufficient memory available.

496 Z\_VERSION\_ERROR

The version requested is not compatible with the library version, or the z\_stream size differs from that used by the library.

In addition, the *msg* field of the *stream* may be set to an error message.

## deflateParams

### Name

500           **deflateParams** — set compression parameters

### Synopsis

```
501 #include <zlib.h>
502 int deflateParams(z_streamp stream, int level, int strategy);
```

### Description

503       The `deflateParams()` function shall dynamically alter the compression parameters  
 504       for the compression stream object `stream`. On entry, `stream` shall refer to a user  
 505       supplied `z_stream` object (a `z_stream_s` structure), already initialized via a call to  
 506       `deflateInit_()` or `deflateInit2_()`.

507       The `level` supplied shall be a value between 0 and 9, or the value  
 508       `Z_DEFAULT_COMPRESSION`. A `level` of 1 requests the highest speed, while a `level` of  
 509       9 requests the highest compression. A `level` of 0 indicates that no compression  
 510       should be used, and the output shall be the same as the input. If the compression  
 511       level is altered by `deflateParams()`, and some data has already been compressed  
 512       with this `stream` (i.e. `total_in` is not zero), and the new `level` requires a different  
 513       underlying compression method, then `stream` shall be flushed by a call to  
 514       `deflate()`.

515       The `strategy` parameter selects the compression strategy to use:

516       `Z_DEFAULT_STRATEGY`

517           use the system default compression strategy. `Z_DEFAULT_STRATEGY` is  
 518           particularly appropriate for text data.

519       `Z_FILTERED`

520           use a compression strategy tuned for data consisting largely of small values  
 521           with a fairly random distribution. `Z_FILTERED` uses more Huffman encoding  
 522           and less string matching than `Z_DEFAULT_STRATEGY`.

523       `Z_HUFFMAN_ONLY`

524           force Huffman encoding only, with no string match.

### Return Value

525       On success, the `deflateParams()` function shall return `z_OK`. Otherwise,  
 526       `deflateParams()` shall return a value as described below to indicate the error.

### Errors

527       On error, `deflateParams()` shall return one of the following error indicators:

528       `Z_STREAM_ERROR`

529           Invalid parameter.

530       `Z_MEM_ERROR`

531           Insufficient memory available.

532            Z\_BUF\_ERROR  
 533            Insufficient space in *stream* to flush the current output.  
 534            In addition, the *msg* field of the *strm* may be set to an error message.

## Application Usage (Informative)

535            Applications should ensure that the *stream* is flushed, e.g. by a call to  
 536            **deflate(stream, Z\_SYNC\_FLUSH)** before calling `deflateParams()`, or ensure that  
 537            there is sufficient space in *next\_out* (as identified by *avail\_out*) to ensure that all  
 538            pending output and all uncompressed input can be flushed in a single call to  
 539            `deflate()`.

540            **Rationale:** Although the `deflateParams()` function should flush pending output and  
 541            compress all pending input, the result is unspecified if there is insufficient space in the  
 542            output buffer. Applications should only call `deflateParams()` when the *stream* is  
 543            effectively empty (flushed).

544            The `deflateParams()` can be used to switch between compression and straight copy of  
 545            the input data, or to switch to a different kind of input data requiring a different strategy.

## deflateReset

### Name

546            `deflateReset — reset compression stream state`

### Synopsis

```
547 #include <zlib.h>
548 int deflateReset(z_streamp stream);
```

### Description

549            The `deflateReset()` function shall reset all state associated with *stream*. All  
 550            pending output shall be discarded, and the counts of processed bytes (*total\_in* and  
 551            *total\_out*) shall be reset to zero.

### Return Value

552            On success, `deflateReset()` shall return `Z_OK`. Otherwise it shall return  
 553            `Z_STREAM_ERROR` to indicate the error.

### Errors

554            On error, `deflateReset()` shall return `Z_STREAM_ERROR`. The following  
 555            conditions shall be treated as an error:

- 556            • The state in *stream* is inconsistent or inappropriate.
- 557            • *stream* is `NULL`.

## deflateSetDictionary

### Name

558      deflateSetDictionary — initialize compression dictionary

### Synopsis

```
559 #include <zlib.h>
560 int deflateSetDictionary(z_streamp stream, const Bytef * dictionary, uInt
561 dictlen);
```

### Description

562      The deflateSetDictionary( ) function shall initialize the compression dictionary  
 563      associated with *stream* using the *dictlen* bytes referenced by *dictionary*.

564      The implementation may silently use a subset of the provided dictionary if the  
 565      dictionary cannot fit in the current window associated with *stream* (see  
 566      deflateInit2\_( )). The application should ensure that the dictionary is sorted such  
 567      that the most commonly used strings occur at the end of the dictionary.

568      If the dictionary is successfully set, the Adler32 checksum of the entire provided  
 569      dictionary shall be stored in the *adler* member of *stream*. This value may be used  
 570      by the decompression system to select the correct dictionary. The compression and  
 571      decompression systems must use the same dictionary.

572      *stream* shall reference an initialized compression stream, with *total\_in* zero (i.e.  
 573      no data has been compressed since the stream was initialized).

### Return Value

574      On success, deflateSetDictionary( ) shall return Z\_OK. Otherwise it shall return  
 575      Z\_STREAM\_ERROR to indicate an error.

### Errors

576      On error, deflateSetDictionary( ) shall return a value as described below:

577      Z\_STREAM\_ERROR

578      The state in *stream* is inconsistent, or *stream* was NULL.

### Application Usage (informative)

579      The application should provide a dictionary consisting of strings {{ed note: do we  
 580      really mean "strings"? Null terminated?}}} that are likely to be encountered in the  
 581      data to be compressed. The application should ensure that the dictionary is sorted  
 582      such that the most commonly used strings occur at the end of the dictionary.

583      The use of a dictionary is optional; however if the data to be compressed is relatively  
 584      short and has a predictable structure, the use of a dictionary can substantially  
 585      improve the compression ratio.

## get\_crc\_table

### Name

586        `get_crc_table` — generate a table for crc calculations

### Synopsis

587        `#include <zlib.h>`  
588        `const uLongf * get_crc_table(void);`

### Description

589        Generate tables for a byte-wise 32-bit CRC calculation based on the polynomial:  
590         $x^{32} + x^{26} + x^{23} + x^{22} + x^{16} + x^{12} + x^{11} + x^{10} + x^8 + x^7 + x^5 + x^4 + x^2 + x + 1$

591        In a multi-threaded application, `get_crc_table()` should be called by one thread to  
592        initialize the tables before any other thread calls any `libz` function.

### Return Value

593        The `get_crc_table()` function shall return a pointer to the first of a set of tables  
594        used internally to calculate CRC-32 values (see `crc32()`).

### Errors

595        None defined.

## gzclose

### Name

596        `gzclose` — close a compressed file stream

### Synopsis

597        `#include <zlib.h>`  
 598        `int gzclose (gzFile file );`

### Description

599        The `gzclose()` function shall close the compressed file stream *file*. If *file* was  
 600        open for writing, `gzclose()` shall first flush any pending output. Any state  
 601        information allocated shall be freed.

### Return Value

602        On success, `gzclose()` shall return `Z_OK`. Otherwise, `gzclose()` shall return an  
 603        error value as described below.

### Errors

604        On error, `gzclose()` may set the global variable `errno` to indicate the error. The  
 605        `gzclose()` shall return a value other than `Z_OK` on error.

606        `Z_STREAM_ERROR`  
 607        *file* was `NULL` (or `Z_NULL`), or did not refer to an open compressed file stream.

608        `Z_ERRNO`  
 609        An error occurred in the underlying base libraries, and the application should  
 610        check `errno` for further information.

611        `Z_BUF_ERROR`  
 612        no compression progress is possible during buffer flush (see `deflate()`).

## gzdopen

### Name

613 gzdopen — open a compressed file

### Synopsis

614 #include <zlib.h>  
615 gzFile gzdopen ( int *fd*, const char \**mode* );

### Description

616 The gzdopen() function shall attempt to associate the open file referenced by *fd*  
617 with a gzFile object. The *mode* argument is based on that of fopen(), but the *mode*  
618 parameter may also contain the following characters:

619 *digit*  
620       set the compression level to *digit*. A low value (e.g. 1) means high speed, while  
621       a high value (e.g. 9) means high compression. A compression level of 0 (zero)  
622       means no compression. See defaultInit2\_() for further details.

623 [*fhR*]  
624       set the compression strategy to [*fhR*]. The letter *f* corresponds to filtered data,  
625       the letter *h* corresponds to Huffman only compression, and the letter *R*  
626       corresponds to Run Length Encoding. See defaultInit2\_() for further details.

627 If *fd* refers to an uncompressed file, and *mode* refers to a read mode, gzdopen() shall  
628 attempt to open the file and return a gzFile object suitable for reading directly from  
629 the file without any decompression.

630 If *mode* is NULL, or if *mode* does not contain one of r, w, or a, gzdopen() shall return  
631 Z\_NULL, and need not set any other error condition.

### Example

632 gzdopen(fileno(stdin), "r");

633 Attempt to associate the standard input with a gzFile object.

### Return Value

634 On success, gzdopen() shall return a gzFile object. On failure, gzdopen() shall  
635 return Z\_NULL and may set errno accordingly.

636 **Note:** At version 1.2.2, zlib does not set errno for several error conditions. Applications  
637 may not be able to determine the cause of an error.

### Errors

638 On error, gzdopen() may set the global variable errno to indicate the error.

## gzeof

### Name

639 gzeof — check for end-of-file on a compressed file stream

### Synopsis

640 #include <zlib.h>  
641 int gzeof (gzFile file );

### Description

642 The `gzeof()` function shall test the compressed file stream *file* for end of file.

### Return Value

643 If *file* was open for reading and end of file has been reached, `gzeof()` shall return 1.  
644 Otherwise, `gzeof()` shall return 0.

### Errors

645 None defined.

**gzerror****Name**

646        `gzerror` — decode an error on a compressed file stream

**Synopsis**

647        `#include <zlib.h>`  
 648        `const char * gzerror (gzFile file, int * errnum);`

**Description**

649        The `gzerror()` function shall return a string describing the last error to have  
 650        occurred associated with the open compressed file stream referred to by `file`. It  
 651        shall also set the location referenced by `errnum` to an integer value that further  
 652        identifies the error.

**Return Value**

653        The `gzerror()` function shall return a string that describes the last error associated  
 654        with the given `file` compressed file stream. This string shall have the format  
 655        "`%s : %s`", with the name of the file, followed by a colon, a space, and the description  
 656        of the error. If the compressed file stream was opened by a call to `gzdopen()`, the  
 657        format of the filename is unspecified.

658        **Rationale:** Although in all current implementations of libz file descriptors are named  
 659        "`<fd : %d>`", the code suggests that this is for debugging purposes only, and may change  
 660        in a future release.

661        It is unspecified if the string returned is determined by the setting of the  
 662        `LC_MESSAGES` category in the current locale.

**Errors**

663        None defined.

## gzflush

### Name

664        `gzflush — flush a compressed file stream`

### Synopsis

665        `#include <zlib.h>`  
 666        `int gzflush(gzFile file, int flush);`

### Description

667        The `gzflush()` function shall flush pending output to the compressed file stream  
 668        identified by `file`, which must be open for writing.

### Flush Operation

670        The parameter `flush` determines which compressed bits are added to the output file.  
 671        If `flush` is `Z_NO_FLUSH`, `gzflush()` may return with some data pending output, and  
 672        not yet written to the file.

673        If `flush` is `Z_SYNC_FLUSH`, `gzflush()` shall flush all pending output to `file` and  
 674        align the output to a byte boundary. There may still be data pending compression  
 675        that is not flushed.

676        If `flush` is `Z_FULL_FLUSH`, all output shall be flushed, as for `Z_SYNC_FLUSH`, and the  
 677        compression state shall be reset. There may still be data pending compression that is  
 678        not flushed.

679        **Rationale:** `Z_SYNC_FLUSH` is intended to ensure that the compressed data contains all the  
 680        data compressed so far, and allows a decompressor to reconstruct all of the input data.  
 681        `Z_FULL_FLUSH` allows decompression to restart from this point if the previous  
 682        compressed data has been lost or damaged. Flushing is likely to degrade the  
 683        performance of the compression system, and should only be used where necessary.

684        If `flush` is set to `Z_FINISH`, all pending uncompressed data shall be compressed and  
 685        all output shall be flushed.

### Return Value

686        On success, `gzflush()` shall return the value `Z_OK`. Otherwise `gzflush()` shall  
 687        return a value to indicate the error, and may set the error number associated with  
 688        the compressed file stream `file`.

689        **Note:** If `flush` is set to `Z_FINISH` and the flush operation is successful, `gzflush()` will  
 690        return `Z_OK`, but the compressed file stream error value may be set to `Z_STREAM_END`.

### Errors

691        On error, `gzwrite()` shall return an error value, and may set the error number  
 692        associated with the stream identified by `file` to indicate the error. Applications may  
 693        use `gzerror()` to access this error value.

694        `Z_ERRNO`

695        An underlying base library function has indicated an error. The global variable  
 696        `errno` may be examined for further information.

697        `Z_STREAM_ERROR`

698                   The stream is invalid, is not open for writing, or is in an invalid state.  
699                   Z\_BUF\_ERROR  
700                   no compression progress is possible (see `deflate()`).  
701                   Z\_MEM\_ERROR  
702                   Insufficient memory available to compress.

## gzgetc

### Name

703        `gzgetc` — read a character from a compressed file

### Synopsis

704        `#include <zlib.h>`  
705        `int gzgetc (gzFile file);`

### Description

706        The `gzgetc()` function shall read the next single character from the compressed file  
707        stream referenced by `file`, which shall have been opened in a read mode (see  
708        `gzopen()` and `gzdopen()`).

### Return Value

709        On success, `gzgetc()` shall return the uncompressed character read, otherwise, on  
710        end of file or error, `gzgetc()` shall return -1.

### Errors

711        On end of file or error, `gzgetc()` shall return -1. Further information can be found  
712        by calling `gzerror()` with a pointer to the compressed file stream.

## gzgets

### Name

713 `gzgets` — read a string from a compressed file

### Synopsis

714 `#include <zlib.h>`  
 715   `char * gzgets (gzFile file, char * buf, int len);`

### Description

716 The `gzgets()` function shall attempt to read data from the compressed file stream  
 717 *file*, uncompressing it into *buf* until either *len*-1 bytes have been inserted into *buf*,  
 718 or until a newline character has been uncompressed into *buf*. A null byte shall be  
 719 appended to the uncompressed data. The *file* shall have been opened in for  
 720 reading (see `gzopen()` and `gzdopen()`).

### Return Value

721 On success, `gzgets()` shall return a pointer to *buf*. Otherwise, `gzgets()` shall  
 722 return `Z_NULL`. Applications may examine the cause using `gzerror()`.

### Errors

723 On error, `gzgets()` shall return `Z_NULL`. The following conditions shall always be  
 724 treated as an error:

*file* is `NULL`, or does not refer to a file open for reading;  
*buf* is `NULL`;  
*len* is less than or equal to zero.

## gzopen

### Name

726        `gzopen` — open a compressed file

### Synopsis

727        `#include <zlib.h>`  
 728        `gzFile gzopen (const char *path , const char *mode );`

### Description

729        The `gzopen()` function shall open the compressed file named by *path*. The *mode*  
 730        argument is based on that of `fopen()`, but the *mode* parameter may also contain the  
 731        following characters:

732        *digit*  
 733                set the compression level to *digit*. A low value (e.g. 1) means high speed, while  
 734                a high value (e.g. 9) means high compression. A compression level of 0 (zero)  
 735                means no compression. See `defaultInit2_()` for further details.

736        *[fhR]*  
 737                set the compression strategy to *[fhR]*. The letter *f* corresponds to filtered data,  
 738                the letter *h* corresponds to Huffman only compression, and the letter *R*  
 739                corresponds to Run Length Encoding. See `defaultInit2_()` for further details.

740        If *path* refers to an uncompressed file, and *mode* refers to a read mode, `gzopen()`  
 741        shall attempt to open the file and return a `gzFile` object suitable for reading directly  
 742        from the file without any decompression.

743        If *path* or *mode* is `NULL`, or if *mode* does not contain one of *r*, *w*, or *a*, `gzopen()` shall  
 744        return `Z_NULL`, and need not set any other error condition.

745        The `gzFile` object is also referred to as a compressed file stream.

### Example

746        `gzopen( "file.gz" , "w6h" );`

747        Attempt to create a new compressed file, `file.gz`, at compression level 6 using  
 748        Huffman only compression.

### Return Value

749        On success, `gzopen()` shall return a `gzFile` object (also known as a *compressed file*  
 750        *stream*). On failure, `gzopen()` shall return `Z_NULL` and may set `errno` accordingly.

751        **Note:** At version 1.2.2, `zlib` does not set `errno` for several error conditions. Applications  
 752        may not be able to determine the cause of an error.

### Errors

753        On error, `gzopen()` may set the global variable `errno` to indicate the error.

## gzprintf

### Name

754        `gzprintf` — format data and compress

### Synopsis

```
755 #include <zlib.h>
756     int gzprintf (gzFile file, const char * fmt, ...);
```

### Description

757        The `gzprintf()` function shall format data as for `fprintf()`, and write the resulting  
758        string to the compressed file stream `file`.

### Return Value

759        The `gzprintf()` function shall return the number of uncompressed bytes actually  
760        written, or a value less than or equal to 0 in the event of an error.

### Errors

761        If `file` is NULL, or refers to a compressed file stream that has not been opened for  
762        writing, `gzprintf()` shall return `Z_STREAM_ERROR`. Otherwise, errors are as for  
763        `gzwrite()`.

## gzputc

### Name

764        `gzputc` — write character to a compressed file

### Synopsis

```
765 #include <zlib.h>
766     int gzputc (gzFile file, int c);
```

### Description

767        The `gzputc()` function shall write the single character `c`, converted from integer to  
768        unsigned character, to the compressed file referenced by `file`, which shall have  
769        been opened in a write mode (see `gzopen()` and `gzdopen()`).

### Return Value

770        On success, `gzputc()` shall return the value written, otherwise `gzputc()` shall  
771        return -1.

### Errors

772        On error, `gzputc()` shall return -1.

## gzputs

### Name

773        `gzputs — string write to a compressed file`

### Synopsis

774        `#include <zlib.h>`  
 775        `int gzputs (gzFile file, const char * s);`

### Description

776        The `gzputs()` function shall write the null terminated string *s* to the compressed file  
 777        referenced by *file*, which shall have been opened in a write mode (see `gzopen()`  
 778        and `gzdopen()`). The terminating null character shall not be written. The `gzputs()`  
 779        function shall return the number of uncompressed bytes actually written.

### Return Value

780        On success, `gzputs()` shall return the number of uncompressed bytes actually  
 781        written to *file*. On error `gzputs()` shall return a value less than or equal to 0.  
 782        Applications may examine the cause using `gzerror()`.

### Errors

783        On error, `gzputs()` shall set the error number associated with the stream identified  
 784        by *file* to indicate the error. Applications should use `gzerror()` to access this error  
 785        value. If *file* is NULL, `gzputs()` shall return `Z_STREAM_ERR`.

786        `Z_ERRNO`

787            An underlying base library function has indicated an error. The global variable  
 788            `errno` may be examined for further information.

789        `Z_STREAM_ERROR`

790            The stream is invalid, is not open for writing, or is in an invalid state.

791        `Z_BUF_ERROR`

792            no compression progress is possible (see `deflate()`).

793        `Z_MEM_ERROR`

794            Insufficient memory available to compress.

## gzread

### Name

795        `gzread — read from a compressed file`

### Synopsis

796        `#include <zlib.h>`  
 797        `int gzread (gzFile file, voidp buf, unsigned int len);`

### Description

798        The `gzread()` function shall read data from the compressed file referenced by `file`,  
 799        which shall have been opened in a read mode (see `gzopen()` and `gzdopen()`). The  
 800        `gzread()` function shall read data from `file`, and uncompress it into `buf`. At most,  
 801        `len` bytes of uncompressed data shall be copied to `buf`. If the file is not compressed,  
 802        `gzread()` shall simply copy data from `file` to `buf` without alteration.

### Return Value

803        On success, `gzread()` shall return the number of bytes decompressed into `buf`. If  
 804        `gzread()` returns 0, either the end-of-file has been reached or an underlying read  
 805        error has occurred. Applications should use `gzerror()` or `gzeof()` to determine  
 806        which occurred. On other errors, `gzread()` shall return a value less than 0 and applications  
 807        may examine the cause using `gzerror()`.

### Errors

808        On error, `gzread()` shall set the error number associated with the stream identified  
 809        by `file` to indicate the error. Applications should use `gzerror()` to access this error  
 810        value.

811        `Z_ERRNO`

812        An underlying base library function has indicated an error. The global variable  
 813        `errno` may be examined for further information.

814        `Z_STREAM_END`

815        End of file has been reached on input.

816        `Z_DATA_ERROR`

817        A CRC error occurred when reading data; the file is corrupt.

818        `Z_STREAM_ERROR`

819        The stream is invalid, or is in an invalid state.

820        `Z_NEED_DICT`

821        A dictionary is needed (see `inflateSetDictionary()`).

822        `Z_MEM_ERROR`

823        Insufficient memory available to decompress.

## gzrewind

### Name

824        `gzrewind` — reset the file-position indicator on a compressed file stream

### Synopsis

825        `#include <zlib.h>`  
826        `int gzrewind(gzFile file);`

### Description

827        The `gzrewind()` function shall set the starting position for the next read on  
828        compressed file stream *file* to the beginning of file. *file* must be open for reading.

829        `gzrewind()` is equivalent to

830        `(int)gzseek(file, 0L, SEEK_SET)`  
831        .

### Return Value

832        On success, `gzrewind()` shall return 0. On error, `gzrewind()` shall return -1, and  
833        may set the error value for *file* accordingly.

### Errors

834        On error, `gzrewind()` shall return -1, indicating that *file* is NULL, or does not  
835        represent an open compressed file stream, or represents a compressed file stream  
836        that is open for writing and is not currently at the beginning of file.

## gzseek

### Name

837        `gzseek` — reposition a file-position indicator in a compressed file stream

### Synopsis

838        `#include <zlib.h>`  
 839        `z_off_t gzseek(gzFile file, z_off_t offset, int whence);`

### Description

840        The `gzseek()` function shall set the file-position indicator for the compressed file  
 841        stream *file*. The file-position indicator controls where the next read or write  
 842        operation on the compressed file stream shall take place. The *offset* indicates a byte  
 843        offset in the uncompressed data. The *whence* parameter may be one of:

844            `SEEK_SET`

845                  the offset is relative to the start of the uncompressed data.

846            `SEEK_CUR`

847                  the offset is relative to the current position in the uncompressed data.

848        **Note:** The value `SEEK_END` need not be supported.

849        If the *file* is open for writing, the new offset must be greater than or equal to the  
 850        current offset. In this case, `gzseek()` shall compress a sequence of null bytes to fill  
 851        the gap from the previous offset to the new offset.

### Return Value

852        On success, `gzseek()` shall return the resulting offset in the file expressed as a byte  
 853        position in the *uncompressed* data stream. On error, `gzseek()` shall return -1, and  
 854        may set the error value for *file* accordingly.

### Errors

855        On error, `gzseek()` shall return -1. The following conditions shall always result in  
 856        an error:

- 857        • *file* is NULL
- 858        • *file* does not represent an open compressed file stream.
- 859        • *file* refers to a compressed file stream that is open for writing, and the newly  
       860        computed offset is less than the current offset.
- 861        • The newly computed offset is less than zero.
- 862        • *whence* is not one of the supported values.

### Application Usage (informative)

863        If *file* is open for reading, the implementation may still need to uncompress all of  
 864        the data up to the new offset. As a result, `gzseek()` may be extremely slow in some  
 865        circumstances.

## gzsetparams

### Name

866        `gzsetparams` — dynamically set compression parameters

### Synopsis

867        `#include <zlib.h>`  
 868        `int gzsetparams (gzFile file, int level, int strategy);`

### Description

869        The `gzsetparams()` function shall set the compression level and compression  
 870        strategy on the compressed file stream referenced by `file`. The compressed file  
 871        stream shall have been opened in a write mode. The `level` and `strategy` are as  
 872        defined in `deflateInit2_`. If there is any data pending writing, it shall be flushed  
 873        before the parameters are updated.

### Return Value

874        On success, the `gzsetparams()` function shall return `Z_OK`.

### Errors

875        On error, `gzsetparams()` shall return one of the following error indications:

876        `Z_STREAM_ERROR`

877            Invalid parameter, or `file` not open for writing.

878        `Z_BUF_ERROR`

879            An internal inconsistency was detected while flushing the previous buffer.

**gztell****Name**

880        `gztell — find position on a compressed file stream`

**Synopsis**

881        `#include <zlib.h>`  
 882        `z_off_t gztell (gzFile file );`

**Description**

883        The `gztell()` function shall return the starting position for the next read or write  
 884        operation on compressed file stream *file*. This position represents the number of  
 885        bytes from the beginning of file in the uncompressed data.

886        `gztell()` is equivalent to

887        `gzseek(file, 0L, SEEK_SET)`

888        .

**Return Value**

889        `gztell()` shall return the current offset in the file expressed as a byte position in the  
 890        *uncompressed* data stream. On error, `gztell()` shall return -1, and may set the error  
 891        value for *file* accordingly.

**Errors**

892        On error, `gztell()` shall return -1, indicating that *file* is NULL, or does not  
 893        represent an open compressed file stream.

## gzwrite

### Name

894        `gzwrite — write to a compressed file`

### Synopsis

895        `#include <zlib.h>`  
 896        `int gzwrite (gzFile file, voidpc buf, unsigned int len);`

### Description

897        The `gzwrite()` function shall write data to the compressed file referenced by `file`,  
 898        which shall have been opened in a write mode (see `gzopen()` and `gzdopen()`). On  
 899        entry, `buf` shall point to a buffer containing `len`bytes of uncompressed data. The  
 900        `gzwrite()` function shall compress this data and write it to `file`. The `gzwrite()`  
 901        function shall return the number of uncompressed bytes actually written.

### Return Value

902        On success, `gzwrite()` shall return the number of uncompressed bytes actually  
 903        written to `file`. On error `gzwrite()` shall return a value less than or equal to 0.  
 904        Applications may examine the cause using `gzerror()`.

### Errors

905        On error, `gzwrite()` shall set the error number associated with the stream identified  
 906        by `file` to indicate the error. Applications should use `gzerror()` to access this error  
 907        value.

908        `Z_ERRNO`

909            An underlying base library function has indicated an error. The global variable  
 910            `errno` may be examined for further information.

911        `Z_STREAM_ERROR`

912            The stream is invalid, is not open for writing, or is in an invalid state.

913        `Z_BUF_ERROR`

914            no compression progress is possible (see `deflate()`).

915        `Z_MEM_ERROR`

916            Insufficient memory available to compress.

## inflate

### Name

917        `inflate` — decompress data

### Synopsis

918        `#include <zlib.h>`  
 919        `int inflate(z_streamp stream, int flush);`

### Description

920        The `inflate()` function shall attempt to decompress data until either the input  
 921        buffer is empty or the output buffer is full. The `stream` references a `z_stream`  
 922        structure. Before the first call to `inflate()`, this structure should have been  
 923        initialized by a call to `inflateInit2_()`.

924        **Note:** `inflateInit2_()` is only in the binary standard; source level applications should  
 925        initialize `stream` via a call to `inflateInit()` or `inflateInit2()`.

926        In addition, the `stream` input and output buffers should have been initialized as  
 927        follows:

928        `next_in`  
 929              should point to the data to be decompressed.  
 930        `avail_in`  
 931              should contain the number of bytes of data in the buffer referenced by `next_in`.  
 932        `next_out`  
 933              should point to a buffer where decompressed data may be placed.  
 934        `avail_out`  
 935              should contain the size in bytes of the buffer referenced by `next_out`

936        The `inflate()` function shall perform one or both of the following actions:

- 937        1. Decompress input data from `next_in` and update `next_in`, `avail_in` and  
        938        `total_in` to reflect the data that has been decompressed.
- 939        2. Fill the output buffer referenced by `next_out`, and update `next_out`,  
        940        `avail_out`, and `total_out` to reflect the decompressed data that has been  
        941        placed there. If `flush` is not `Z_NO_FLUSH`, and `avail_out` indicates that there is  
        942        still space in output buffer, this action shall always occur (see below for further  
        943        details).

944        The `inflate()` function shall return when either `avail_in` reaches zero (indicating  
 945        that all the input data has been compressed), or `avail_out` reaches zero (indicating  
 946        that the output buffer is full).

947        On success, the `inflate()` function shall set the `adler` field of the `stream` to the  
 948        Adler-32 checksum of all the input data compressed so far (represented by  
 949        `total_in`).

### Flush Operation

951       The parameter *flush* determines when uncompressed bytes are added to the output  
 952       buffer in *next\_out*. If *flush* is `Z_NO_FLUSH`, `inflate()` may return with some data  
 953       pending output, and not yet added to the output buffer.  
 954       If *flush* is `Z_SYNC_FLUSH`, `inflate()` shall flush all pending output to *next\_out*,  
 955       and update *next\_out* and *avail\_out* accordingly.  
 956       If *flush* is set to `Z_BLOCK`, `inflate()` shall stop adding data to the output buffer if  
 957       and when the next compressed block boundary is reached (see RFC 1951: DEFLATE  
 958       Compressed Data Format Specification).  
 959       If *flush* is set to `Z_FINISH`, all of the compressed input shall be decompressed and  
 960       added to the output. If there is insufficient output space (i.e. the compressed input  
 961       data uncompresses to more than *avail\_out* bytes), then `inflate()` shall fail and  
 962       return `Z_BUF_ERROR`.

## Return Value

963       On success, `inflate()` shall return `Z_OK` if decompression progress has been made,  
 964       or `Z_STREAM_END` if all of the input data has been decompressed and there was  
 965       sufficient space in the output buffer to store the uncompressed result. On error,  
 966       `inflate()` shall return a value to indicate the error.

967       **Note:** If `inflate()` returns `Z_OK` and has set *avail\_out* to zero, the function should be  
 968       called again with the same value for *flush*, and with updated *next\_out* and *avail\_out*  
 969       until `inflate()` returns with either `Z_OK` or `Z_STREAM_END` and a non-zero  
 970       *avail\_out*.

971       On success, `inflate()` shall set the *adler* to the Adler-32 checksum of the output  
 972       produced so far (i.e. *total\_out* bytes).

## Errors

973       On error, `inflate()` shall return a value as described below, and may set the *msg*  
 974       field of *stream* to point to a string describing the error:

975       `Z_BUF_ERROR`

976       No progress is possible; either *avail\_in* or *avail\_out* was zero.

977       `Z_MEM_ERROR`

978       Insufficient memory.

979       `Z_STREAM_ERROR`

980       The state (as represented in *stream*) is inconsistent, or *stream* was `NULL`.

981       `Z_NEED_DICT`

982       A preset dictionary is required. The *adler* field shall be set to the Adler-32  
 983       checksum of the dictionary chosen by the compressor.

**inflateEnd****Name**

984           **inflateEnd** — free decompression stream state

**Synopsis**

985       **#include <zlib.h>**  
 986       **int inflateEnd(z\_streamp stream);**

**Description**

987       The **inflateEnd()** function shall free all allocated state information referenced by  
 988       *stream*. All pending output is discarded, and unprocessed input is ignored.

**Return Value**

989       On success, **inflateEnd()** shall return **Z\_OK**. Otherwise it shall return  
 990       **Z\_STREAM\_ERROR** to indicate the error.

**Errors**

991       On error, **inflateEnd()** shall return **Z\_STREAM\_ERROR**. The following conditions  
 992       shall be treated as an error:

- 993       • The state in *stream* is inconsistent.  
 994       • *stream* is **NULL**.  
 995       • The *zfree* function pointer is **NULL**.

**inflateInit2\_****Name**

996           **inflateInit2\_** — initialize decompression system

**Synopsis**

```
997 #include <zlib.h>
998   int inflateInit2_ (z_streamp strm, int windowBits, char * version, int
999   stream_size);
```

**Description**

1000 The **inflateInit2\_()** function shall initialize the decompression system. On entry,  
 1001 *strm* shall refer to a user supplied **z\_stream** object (a **z\_stream\_s** structure). The  
 1002 following fields shall be set on entry:

1003 *zalloc*

1004       a pointer to an *alloc\_func* function, used to allocate state information. If this is  
 1005 NULL, a default allocation function will be used.

1006 *zfree*

1007       a pointer to a *free\_func* function, used to free memory allocated by the *zalloc*  
 1008 function. If this is NULL a default free function will be used.

1009 *opaque*

1010       If *alloc\_func* is not NULL, *opaque* is a user supplied pointer to data that will be  
 1011 passed to the *alloc\_func* and *free\_func* functions.

1012       If the *version* requested is not compatible with the version implemented, or if the  
 1013 size of the **z\_stream\_s** structure provided in *stream\_size* does not match the size  
 1014 in the library implementation, **inflateInit2\_()** shall fail, and return  
 1015 **Z\_VERSION\_ERROR**.

1016       The *windowBits* parameter shall be a base 2 logarithm of the maximum window size  
 1017 to use, and shall be a value between 8 and 15. If the input data was compressed with  
 1018 a larger window size, subsequent attempts to decompress this data will fail with  
 1019 **Z\_DATA\_ERROR**, rather than try to allocate a larger window.

1020       The **inflateInit2\_()** function is not in the source standard; it is only in the binary  
 1021 standard. Source applications should use the **inflateInit2()** macro.

**Return Value**

1022       On success, the **inflateInit2\_()** function shall return **Z\_OK**. Otherwise,  
 1023 **inflateInit2\_()** shall return a value as described below to indicate the error.

**Errors**

1024       On error, **inflateInit2\_()** shall return one of the following error indicators:

1025       **Z\_STREAM\_ERROR**

1026       Invalid parameter.

1027       **Z\_MEM\_ERROR**

1028       Insufficient memory available.

1029 Z\_VERSION\_ERROR

1030       The version requested is not compatible with the library version, or the  
1031       z\_stream size differs from that used by the library.

1032      In addition, the *msg* field of the *strm* may be set to an error message.

**inflateInit\_****Name**

1033   **inflateInit\_** — initialize decompression system

**Synopsis**

```
1034 #include <zlib.h>
1035 int inflateInit_(z_streamp stream, const char * version, int stream_size);
```

**Description**

1036 The **inflateInit\_()** function shall initialize the decompression system. On entry,  
 1037 *stream* shall refer to a user supplied *z\_stream* object (a *z\_stream\_s* structure). The  
 1038 following fields shall be set on entry:

1039 *zalloc*

1040   a pointer to an *alloc\_func* function, used to allocate state information. If this is  
 1041   NULL, a default allocation function will be used.

1042 *zfree*

1043   a pointer to a *free\_func* function, used to free memory allocated by the *zalloc*  
 1044   function. If this is NULL a default free function will be used.

1045 *opaque*

1046   If *alloc\_func* is not NULL, *opaque* is a user supplied pointer to data that will be  
 1047   passed to the *alloc\_func* and *free\_func* functions.

1048 If the *version* requested is not compatible with the version implemented, or if the  
 1049 size of the *z\_stream\_s* structure provided in *stream\_size* does not match the size  
 1050 in the library implementation, **inflateInit\_()** shall fail, and return  
 1051 *Z\_VERSION\_ERROR*.

1052 The **inflateInit\_()** function is not in the source standard; it is only in the binary  
 1053 standard. Source applications should use the **inflateInit()** macro.

1054 The **inflateInit\_()** shall be equivalent to

```
1055 inflateInit2_(strm, DEF_WBITS, version, stream_size);
```

**Return Value**

1056 On success, the **inflateInit\_()** function shall return *z\_OK*. Otherwise,  
 1057 **inflateInit\_()** shall return a value as described below to indicate the error.

**Errors**

1058 On error, **inflateInit\_()** shall return one of the following error indicators:

1059 *Z\_STREAM\_ERROR*

1060   Invalid parameter.

1061 *Z\_MEM\_ERROR*

1062   Insufficient memory available.

1063 *Z\_VERSION\_ERROR*

1064           The version requested is not compatible with the library version, or the  
 1065           z\_stream size differs from that used by the library.

1066           In addition, the *msg* field of the *strm* may be set to an error message.

## **inflateReset**

### **Name**

1067           **inflateReset** — reset decompression stream state

### **Synopsis**

```
1068 #include <zlib.h>
1069 int inflateReset(z_streamp stream);
```

### **Description**

1070           The *inflateReset()* function shall reset all state associated with *stream*. All  
 1071           pending output shall be discarded, and the counts of processed bytes (*total\_in* and  
 1072           *total\_out*) shall be reset to zero.

### **Return Value**

1073           On success, *inflateReset()* shall return *Z\_OK*. Otherwise it shall return  
 1074           *Z\_STREAM\_ERROR* to indicate the error.

### **Errors**

1075           On error, *inflateReset()* shall return *Z\_STREAM\_ERROR*. The following  
 1076           conditions shall be treated as an error:

- The state in *stream* is inconsistent or inappropriate.
- *stream* is NULL.

## inflateSetDictionary

### Name

1079       **inflateSetDictionary** — initialize decompression dictionary

### Synopsis

```
1080 #include <zlib.h>
1081 int inflateSetDictionary(z_streamp stream, const Bytef * dictionary, uInt
1082 dictlen);
```

### Description

1083       The **inflateSetDictionary()** function shall initialize the decompression  
 1084       dictionary associated with *stream* using the *dictlen* bytes referenced by  
 1085       *dictionary*.

1086       The **inflateSetDictionary()** function should be called immediately after a call to  
 1087       **inflate()** has failed with return value Z\_NEED\_DICT. The *dictionary* must have  
 1088       the same Adler-32 checksum as the dictionary used for the compression (see  
 1089       **deflateSetDictionary()**).

1090       *stream* shall reference an initialized decompression stream, with *total\_in* zero (i.e.  
 1091       no data has been decompressed since the stream was initialized).

### Return Value

1092       On success, **inflateSetDictionary()** shall return Z\_OK. Otherwise it shall return  
 1093       a value as indicated below.

### Errors

1094       On error, **inflateSetDictionary()** shall return a value as described below:

1095       Z\_STREAM\_ERROR

1096           The state in *stream* is inconsistent, or *stream* was NULL.

1097       Z\_DATA\_ERROR

1098           The Adler-32 checksum of the supplied dictionary does not match that used for  
 1099           the compression.

### Application Usage (informative)

1100       The application should provide a dictionary consisting of strings {{ed note: do we  
 1101       really mean "strings"? Null terminated?}}} that are likely to be encountered in the  
 1102       data to be compressed. The application should ensure that the dictionary is sorted  
 1103       such that the most commonly used strings occur at the end of the dictionary.

1104       The use of a dictionary is optional; however if the data to be compressed is relatively  
 1105       short and has a predictable structure, the use of a dictionary can substantially  
 1106       improve the compression ratio.

## inflateSync

### Name

1107        inflateSync — advance compression stream to next sync point

### Synopsis

1108        #include <zlib.h>  
1109        int inflateSync(z\_streamp stream);

### Description

1110        The `inflateSync()` function shall advance through the compressed data in *stream*,  
1111        skipping any invalid compressed data, until the next full flush point is reached, or  
1112        all input is exhausted. See the description for `deflate()` with flush level  
1113        `Z_FULL_FLUSH`. No output is placed in *next\_out*.

### Return Value

1114        On success, `inflateSync()` shall return `Z_OK`, and update the *next\_in*, *avail\_in*,  
1115        and, *total\_in* fields of *stream* to reflect the number of bytes of compressed data  
1116        that have been skipped. Otherwise, `inflateSync()` shall return a value as described  
1117        below to indicate the error.

### Errors

1118        On error, `inflateSync()` shall return a value as described below:

1119        `Z_STREAM_ERROR`

1120        The state (as represented in *stream*) is inconsistent, or *stream* was `NULL`.

1121        `Z_BUF_ERROR`

1122        There is no data available to skip over.

1123        `Z_DATA_ERROR`

1124        No sync point was found.

## inflateSyncPoint

### Name

1125 inflateSyncPoint — test for synchronization point

### Synopsis

1126 #include <zlib.h>  
1127 int inflateSyncPoint(z\_streamp stream);

### Description

1128 The `inflateSyncPoint()` function shall return a non-zero value if the compressed  
1129 data stream referenced by `stream` is at a synchronization point.

### Return Value

1130 If the compressed data in `stream` is at a synchronization point (see `deflate()` with a  
1131 flush level of `Z_SYNC_FLUSH` or `Z_FULL_FLUSH`), `inflateSyncPoint()` shall return a  
1132 non-zero value, other than `Z_STREAM_ERROR`. Otherwise, if the `stream` is valid,  
1133 `inflateSyncPoint()` shall return 0. If `stream` is invalid, or in an invalid state,  
1134 `inflateSyncPoint()` shall return `Z_STREAM_ERROR` to indicate the error.

### Errors

1135 On error, `inflateSyncPoint()` shall return a value as described below:

1136 `Z_STREAM_ERROR`

1137 The state (as represented in `stream`) is inconsistent, or `stream` was NULL.

## uncompress

### Name

1138    uncompress — uncompress data

### Synopsis

```
1139 #include <zlib.h>
1140 int uncompress(Bytef * dest, uLongf * destLen, const Bytef * source, uLong
1141 sourceLen);
```

### Description

1142 The `uncompress()` function shall attempt to uncompress `sourceLen` bytes of data in  
 1143 the buffer `source`, placing the result in the buffer `dest`.

1144 On entry, `destLen` should point to a value describing the size of the `dest` buffer. The  
 1145 application should ensure that this value is large enough to hold the entire  
 1146 uncompressed data.

1147 **Note:** The LSB does not describe any mechanism by which a compressor can  
 1148 communicate the size required to the uncompressor.

1149 On successful exit, the variable referenced by `destLen` shall be updated to hold the  
 1150 length of uncompressed data in `dest`.

### Return Value

1151 On success, `uncompress()` shall return `Z_OK`. Otherwise, `uncompress()` shall  
 1152 return a value to indicate the error.

### Errors

1153 On error, `uncompress()` shall return a value as described below:

1154 `Z_BUF_ERROR`

1155    The buffer `dest` was not large enough to hold the uncompressed data.

1156 `Z_MEM_ERROR`

1157    Insufficient memory.

1158 `Z_DATA_ERROR`

1159    The compressed data (referenced by `source`) was corrupted.

**zError****Name**

1160        `zError` — translate error number to string

**Synopsis**

```
1161 #include <zlib.h>
1162 const char * zError(int err);
```

**Description**

1163        The `zError()` function shall return the string identifying the error associated with  
 1164        `err`. This allows for conversion from error code to string for functions such as  
 1165        `compress()` and `uncompress()`, that do not always set the string version of an error.

**Return Value**

1166        The `zError()` function shall return a the string identifying the error associated with  
 1167        `err`, or NULL if `err` is not a valid error code.

1168        It is unspecified if the string returned is determined by the setting of the  
 1169        `LC_MESSAGES` category in the current locale.

**Errors**

1170        None defined.

**zlibVersion****Name**

1171        `zlibVersion` — discover library version at run time

**Synopsis**

```
1172 #include <zlib.h>
1173 const char * zlibVersion (void);
```

**Description**

1174        The `zlibVersion()` function shall return the string identifying the interface version  
 1175        at the time the library was built.

1176        Applications should compare the value returned from `zlibVersion()` with the  
 1177        macro constant `ZLIB_VERSION` for compatibility.

**Return Value**

1178        The `zlibVersion()` function shall return a the string identifying the version of the  
 1179        library currently implemented.

**Errors**

1180        None defined.

**14.5 Interfaces for libncurses**

1181        Table 14-3 defines the library name and shared object name for the libncurses library

1182

**Table 14-3 libncurses Definition**

Library:	libncurses
SONAME:	libncurses.so.5

1183

The Parameters or return value of the following interface have had the const qualifier added as shown here.

1186

```
extern const char *keyname (int);
extern int mvscanw (int, int, const char *, ...);
extern int mvwscanw (WINDOW *, int, int, const char *, ...);
extern SCREEN *newterm (const char *, FILE *, FILE *);
extern int scanw (const char *, ...);
extern int vwscanw (WINDOW *, const char *, va_list);
extern int vw_scanw (WINDOW *, const char *, va_list);
extern int wscanw (WINDOW *, const char *, ...);
```

1194

The behavior of the interfaces in this library is specified by the following specifications:

1195

[SUS-CURSES] X/Open Curses

## 14.5.1 Curses

1197

### 14.5.1.1 Interfaces for Curses

1198

An LSB conforming implementation shall provide the generic functions for Curses specified in Table 14-4, with the full mandatory functionality as described in the referenced underlying specification.

1201

**Table 14-4 libncurses - Curses Function Interfaces**

addch [SUS-CURSES]	addchnstr [SUS-CURSES]	addchstr [SUS-CURSES]	addnstr [SUS-CURSES]
addstr [SUS-CURSES]	attr_get [SUS-CURSES]	attr_off [SUS-CURSES]	attr_on [SUS-CURSES]
attr_set [SUS-CURSES]	attroff [SUS-CURSES]	attron [SUS-CURSES]	attrset [SUS-CURSES]
baudrate [SUS-CURSES]	beep [SUS-CURSES]	bkgd [SUS-CURSES]	bkgdset [SUS-CURSES]
border [SUS-CURSES]	box [SUS-CURSES]	can_change_color [SUS-CURSES]	cbreak [SUS-CURSES]
chgat [SUS-CURSES]	clear [SUS-CURSES]	clearok [SUS-CURSES]	clrtobot [SUS-CURSES]
clrtoeol [SUS-CURSES]	color_content [SUS-CURSES]	color_set [SUS-CURSES]	copywin [SUS-CURSES]
curs_set [SUS-CURSES]	def_prog_mode [SUS-CURSES]	def_shell_mode [SUS-CURSES]	del_curterm [SUS-CURSES]
delay_output [SUS-CURSES]	delch [SUS-CURSES]	deleteln [SUS-CURSES]	delscreen [SUS-CURSES]
delwin	derwin	doupdate	dupwin

[SUS-CURSES]	[SUS-CURSES]	[SUS-CURSES]	[SUS-CURSES]
echo [SUS-CURSES]	echochar [SUS-CURSES]	endwin [SUS-CURSES]	erase [SUS-CURSES]
erasechar [SUS-CURSES]	filter [SUS-CURSES]	flash [SUS-CURSES]	flushinp [SUS-CURSES]
getbkgd [SUS-CURSES]	getch [SUS-CURSES]	getnstr [SUS-CURSES]	getstr [SUS-CURSES]
getwin [SUS-CURSES]	halfdelay [SUS-CURSES]	has_colors [SUS-CURSES]	has_ic [SUS-CURSES]
has_il [SUS-CURSES]	hline [SUS-CURSES]	idcok [SUS-CURSES]	idlok [SUS-CURSES]
immedok [SUS-CURSES]	inch [SUS-CURSES]	inchnstr [SUS-CURSES]	inchstr [SUS-CURSES]
init_color [SUS-CURSES]	init_pair [SUS-CURSES]	initscr [SUS-CURSES]	innstr [SUS-CURSES]
insch [SUS-CURSES]	insdelln [SUS-CURSES]	insertln [SUS-CURSES]	insnstr [SUS-CURSES]
insstr [SUS-CURSES]	instr [SUS-CURSES]	intrflush [SUS-CURSES]	is_linetouched [SUS-CURSES]
is_wintouched [SUS-CURSES]	isendwin [SUS-CURSES]	keyname [SUS-CURSES]	keypad [SUS-CURSES]
killchar [SUS-CURSES]	leaveok [SUS-CURSES]	longname [SUS-CURSES]	meta [SUS-CURSES]
move [SUS-CURSES]	mvaddch [SUS-CURSES]	mvaddchnstr [SUS-CURSES]	mvaddchstr [SUS-CURSES]
mvaddnstr [SUS-CURSES]	mvaddstr [SUS-CURSES]	mvchgat [SUS-CURSES]	mvcur [SUS-CURSES]
mvdelch [SUS-CURSES]	mvderwin [SUS-CURSES]	mvgetch [SUS-CURSES]	mvgetnstr [SUS-CURSES]
mvgetstr [SUS-CURSES]	mvhline [SUS-CURSES]	mvinch [SUS-CURSES]	mvinchnstr [SUS-CURSES]
mvinchstr [SUS-CURSES]	mvinnstr [SUS-CURSES]	mvinsch [SUS-CURSES]	mvinsnstr [SUS-CURSES]
mvinsstr [SUS-CURSES]	mvinstr [SUS-CURSES]	mvprintw [SUS-CURSES]	mvscanw [SUS-CURSES]
mvvline [SUS-CURSES]	mvwaddch [SUS-CURSES]	mvwaddchnstr [SUS-CURSES]	mvwaddchstr [SUS-CURSES]
mvwaddnstr [SUS-CURSES]	mvwaddstr [SUS-CURSES]	mvwchgat [SUS-CURSES]	mvwdelch [SUS-CURSES]
mvwgetch [SUS-CURSES]	mvwgetnstr [SUS-CURSES]	mvwgetstr [SUS-CURSES]	mvwhline [SUS-CURSES]

<code>mvwin</code> [SUS-CURSES]	<code>mvwinch</code> [SUS-CURSES]	<code>mvwinchnstr</code> [SUS-CURSES]	<code>mvwinchstr</code> [SUS-CURSES]
<code>mvwinnstr</code> [SUS-CURSES]	<code>mvwinsch</code> [SUS-CURSES]	<code>mvwinsnstr</code> [SUS-CURSES]	<code>mvwinsstr</code> [SUS-CURSES]
<code>mvwinstr</code> [SUS-CURSES]	<code>mvwprintw</code> [SUS-CURSES]	<code>mvwscanw</code> [SUS-CURSES]	<code>mvwvline</code> [SUS-CURSES]
<code>napms</code> [SUS-CURSES]	<code>newpad</code> [SUS-CURSES]	<code>newterm</code> [SUS-CURSES]	<code>newwin</code> [SUS-CURSES]
<code>nl</code> [SUS-CURSES]	<code>nocbreak</code> [SUS-CURSES]	<code>nodelay</code> [SUS-CURSES]	<code>noecho</code> [SUS-CURSES]
<code>nonl</code> [SUS-CURSES]	<code>noqflush</code> [SUS-CURSES]	<code>noraw</code> [SUS-CURSES]	<code>notimeout</code> [SUS-CURSES]
<code>overlay</code> [SUS-CURSES]	<code>overwrite</code> [SUS-CURSES]	<code>pair_content</code> [SUS-CURSES]	<code>pechochar</code> [SUS-CURSES]
<code>pnoutrefresh</code> [SUS-CURSES]	<code>prefresh</code> [SUS-CURSES]	<code>printw</code> [SUS-CURSES]	<code>putp</code> [SUS-CURSES]
<code>putwin</code> [SUS-CURSES]	<code>qiflush</code> [SUS-CURSES]	<code>raw</code> [SUS-CURSES]	<code>redrawwin</code> [SUS-CURSES]
<code>refresh</code> [SUS-CURSES]	<code>reset_prog_mode</code> [SUS-CURSES]	<code>reset_shell_mode</code> [SUS-CURSES]	<code>resetty</code> [SUS-CURSES]
<code>restartterm</code> [SUS-CURSES]	<code>riponline</code> [SUS-CURSES]	<code>savetty</code> [SUS-CURSES]	<code>scanw</code> [SUS-CURSES]
<code>scr_dump</code> [SUS-CURSES]	<code>scr_init</code> [SUS-CURSES]	<code>scr_restore</code> [SUS-CURSES]	<code>scr_set</code> [SUS-CURSES]
<code>scrl</code> [SUS-CURSES]	<code>scroll</code> [SUS-CURSES]	<code>scrollok</code> [SUS-CURSES]	<code>set_curterm</code> [SUS-CURSES]
<code>set_term</code> [SUS-CURSES]	<code>setsqrreg</code> [SUS-CURSES]	<code>setupterm</code> [SUS-CURSES]	<code>slk_attr_set</code> [SUS-CURSES]
<code>slk_attroff</code> [SUS-CURSES]	<code>slk_attron</code> [SUS-CURSES]	<code>slk_attrset</code> [SUS-CURSES]	<code>slk_clear</code> [SUS-CURSES]
<code>slk_color</code> [SUS-CURSES]	<code>slk_init</code> [SUS-CURSES]	<code>slk_label</code> [SUS-CURSES]	<code>slk_noutrefresh</code> [SUS-CURSES]
<code>slk_refresh</code> [SUS-CURSES]	<code>slk_restore</code> [SUS-CURSES]	<code>slk_set</code> [SUS-CURSES]	<code>slk_touch</code> [SUS-CURSES]
<code>standend</code> [SUS-CURSES]	<code>standout</code> [SUS-CURSES]	<code>start_color</code> [SUS-CURSES]	<code>subpad</code> [SUS-CURSES]
<code>subwin</code> [SUS-CURSES]	<code>syncok</code> [SUS-CURSES]	<code>termattrs</code> [SUS-CURSES]	<code>termname</code> [SUS-CURSES]
<code>tgetent</code> [SUS-CURSES]	<code>tgetflag</code> [SUS-CURSES]	<code>tgetnum</code> [SUS-CURSES]	<code>tgetstr</code> [SUS-CURSES]
<code>tgoto</code>	<code>tigetflag</code>	<code>tigetnum</code>	<code>tigetstr</code>

[SUS-CURSES]	[SUS-CURSES]	[SUS-CURSES]	[SUS-CURSES]
timeout [SUS-CURSES]	touchline [SUS-CURSES]	touchwin [SUS-CURSES]	tparm [SUS-CURSES]
tputs [SUS-CURSES]	typeahead [SUS-CURSES]	unctrl [SUS-CURSES]	ungetch [SUS-CURSES]
untouchwin [SUS-CURSES]	use_env [SUS-CURSES]	vidattr [SUS-CURSES]	vidputs [SUS-CURSES]
vline [SUS-CURSES]	vwprintw [SUS-CURSES]	vw_scanw [SUS-CURSES]	vwprintw [SUS-CURSES]
vwscanw [SUS-CURSES]	waddch [SUS-CURSES]	waddchnstr [SUS-CURSES]	waddchstr [SUS-CURSES]
waddnstr [SUS-CURSES]	waddstr [SUS-CURSES]	wattr_get [SUS-CURSES]	wattr_off [SUS-CURSES]
wattr_on [SUS-CURSES]	wattr_set [SUS-CURSES]	wattroff [SUS-CURSES]	wattron [SUS-CURSES]
wattrset [SUS-CURSES]	wbkgd [SUS-CURSES]	wbkgdset [SUS-CURSES]	wborder [SUS-CURSES]
wchgat [SUS-CURSES]	wclear [SUS-CURSES]	wclrbot [SUS-CURSES]	wclrtobot [SUS-CURSES]
wcolor_set [SUS-CURSES]	wcursyncup [SUS-CURSES]	wdelch [SUS-CURSES]	wdeleteln [SUS-CURSES]
wechochar [SUS-CURSES]	werase [SUS-CURSES]	wgetch [SUS-CURSES]	wgetnstr [SUS-CURSES]
wgetstr [SUS-CURSES]	whline [SUS-CURSES]	winch [SUS-CURSES]	winchnstr [SUS-CURSES]
winchstr [SUS-CURSES]	winnstr [SUS-CURSES]	winsch [SUS-CURSES]	winsdelln [SUS-CURSES]
wininsertln [SUS-CURSES]	winsnstr [SUS-CURSES]	winsstr [SUS-CURSES]	winstr [SUS-CURSES]
wmove [SUS-CURSES]	wnoutrefresh [SUS-CURSES]	wprintw [SUS-CURSES]	wredrawln [SUS-CURSES]
wrefresh [SUS-CURSES]	wscanw [SUS-CURSES]	wscrl [SUS-CURSES]	wsetscreg [SUS-CURSES]
wstandend [SUS-CURSES]	wstandout [SUS-CURSES]	wsyncdown [SUS-CURSES]	wsyncup [SUS-CURSES]
wtimeout [SUS-CURSES]	wtouchln [SUS-CURSES]	wvline [SUS-CURSES]	

1202

1203

1204

1205

An LSB conforming implementation shall provide the generic data interfaces for Curses specified in Table 14-5, with the full mandatory functionality as described in the referenced underlying specification.

1206

**Table 14-5 libncurses - Curses Data Interfaces**

1207

COLORS [SUS-CURSES]	COLOR_PAIRS [SUS-CURSES]	COLS [SUS-CURSES]	LINES [SUS-CURSES]
acs_map [SUS-CURSES]	cur_term [SUS-CURSES]	curscr [SUS-CURSES]	stdscr [SUS-CURSES]

## 14.6 Data Definitions for libncurses

1208

This section defines global identifiers and their values that are associated with interfaces contained in libncurses. These definitions are organized into groups that correspond to system headers. This convention is used as a convenience for the reader, and does not imply the existence of these headers, or their content. Where an interface is defined as requiring a particular system header file all of the data definitions for that system header file presented here shall be in effect.

1214

This section gives data definitions to promote binary application portability, not to repeat source interface definitions available elsewhere. System providers and application developers should use this ABI to supplement - not to replace - source interface definition specifications.

1219

This specification uses the ISO C (1999) C Language as the reference programming language, and data definitions are specified in ISO C format. The C language is used here as a convenient notation. Using a C language description of these data objects does not preclude their use by other programming languages.

### 14.6.1 curses.h

1222

```

1223 #define ERR      (-1)
1224 #define OK       (0)
1225 #define ACS_RARROW      (acs_map['+'])
1226 #define ACS_LARROW      (acs_map[' , '])
1227 #define ACS_UARROW      (acs_map['- '])
1228 #define ACS_DARROW      (acs_map['.' ])
1229 #define ACS_BLOCK        (acs_map['0'])
1230 #define ACS_CKBOARD      (acs_map['a'])
1231 #define ACS_DEGREE       (acs_map['f'])
1232 #define ACS_PLMINUS      (acs_map['g'])
1233 #define ACS_BOARD        (acs_map['h'])
1234 #define ACS_LANTERN      (acs_map['i'])
1235 #define ACS_LRCORNER     (acs_map['j'])
1236 #define ACS_URCORNER     (acs_map['k'])
1237 #define ACS_ULCORNER     (acs_map['l'])
1238 #define ACS_LLCORNER     (acs_map['m'])
1239 #define ACS_PLUS          (acs_map['n'])
1240 #define ACS_S1            (acs_map['o'])
1241 #define ACS_HLINE         (acs_map['q'])
1242 #define ACS_S9            (acs_map['s'])
1243 #define ACS_LTEE          (acs_map['t'])
1244 #define ACS_RTEE          (acs_map['u'])
1245 #define ACS_BTEE          (acs_map['v'])
1246 #define ACS_TTEE          (acs_map['w'])
1247 #define ACS_VLINE         (acs_map['x'])
1248 #define ACS_DIAMOND       (acs_map['`'])
1249 #define ACS_BULLET         (acs_map['~'])
1250 #define getmaxyx(win,y,x) \
1251             (y=(win)?((win)->_maxy+1):ERR,x=(win)?((win)->_maxx+1):ERR)
1252 #define getbegyx(win,y,x) \
1253             (y=(win)?(win)->_begy:ERR,x=(win)?(win)->_begx:ERR)

```

```

1254     #define getyx(win,y,x) \
1255         (y=(win)?(win)->_cury:ERR,x=(win)?(win)->_curx:ERR)
1256     #define getparyx(win,y,x) \
1257         (y=(win)?(win)->_parx:ERR,x=(win)?(win)->_parx:ERR)
1258
1259     #define WA_ALTCHARSET A_ALTCHARSET
1260     #define WA_ATTRIBUTES A_ATTRIBUTES
1261     #define WA_BLINK A_BLINK
1262     #define WA_BOLD A_BOLD
1263     #define WA_DIM A_DIM
1264     #define WA_HORIZONTAL A_HORIZONTAL
1265     #define WA_INVIS A_INVIS
1266     #define WA_LEFT A_LEFT
1267     #define WA_LOW A_LOW
1268     #define WA_NORMAL A_NORMAL
1269     #define WA_PROTECT A_PROTECT
1270     #define WA_REVERSE A_REVERSE
1271     #define WA_RIGHT A_RIGHT
1272     #define WA_STANDOUT A_STANDOUT
1273     #define WA_TOP A_TOP
1274     #define WA_UNDERLINE A_UNDERLINE
1275     #define WA_VERTICAL A_VERTICAL
1276     #define A_REVERSE NCURSES_BITS(1UL,10)
1277
1278     #define COLOR_BLACK 0
1279     #define COLOR_RED 1
1280     #define COLOR_GREEN 2
1281     #define COLOR_YELLOW 3
1282     #define COLOR_BLUE 4
1283     #define COLOR_MAGENTA 5
1284     #define COLOR_CYAN 6
1285     #define COLOR_WHITE 7
1286
1287     #define _SUBWIN 0x01
1288     #define _ENDLINE 0x02
1289     #define _FULLWIN 0x04
1290     #define _ISPAD 0x10
1291     #define _HASMOVED 0x20
1292
1293     typedef unsigned char bool;
1294
1295     typedef unsigned long int chtype;
1296     typedef struct screen SCREEN;
1297     typedef struct _win_st WINDOW;
1298     typedef chtype attr_t;
1299     typedef struct {
1300         attr_t attr;
1301         wchar_t chars[5];
1302     } cchar_t;
1303     struct pdat {
1304         short _pad_y;
1305         short _pad_x;
1306         short _pad_top;
1307         short _pad_left;
1308         short _pad_bottom;
1309         short _pad_right;
1310     };
1311
1312     struct _win_st {
1313         short _cury;
1314         short _curx;
1315         short _maxy;
1316         short _maxx;
1317         short _begy;

```

```

1318     short _begx;
1319     short _flags;
1320     attr_t _attrs;
1321     chtype _bkgd;
1322     bool _notimeout;
1323     bool _clear;
1324     bool _leaveok;
1325     bool _scroll;
1326     bool _idlok;
1327     bool _idcok;
1328     bool _immed;
1329     bool _sync;
1330     bool _use_keypad;
1331     int _delay;
1332     struct ldat *_line;
1333     short _regtop;
1334     short _regbottom;
1335     int _parx;
1336     int _pary;
1337     WINDOW *_parent;
1338     struct pdat _pad;
1339     short _yoffset;
1340     cchar_t _bkgrnd;
1341 };
1342
1343 #define KEY_CODE_YES      0400
1344 #define KEY_BREAK        0401
1345 #define KEY_MIN          0401
1346 #define KEY_DOWN         0402
1347 #define KEY_UP           0403
1348 #define KEY_LEFT          0404
1349 #define KEY_RIGHT         0405
1350 #define KEY_HOME          0406
1351 #define KEY_BACKSPACE    0407
1352 #define KEY_F0            0410
1353 #define KEY_DL            0510
1354 #define KEY_IL            0511
1355 #define KEY_DC            0512
1356 #define KEY_IC            0513
1357 #define KEY_EIC           0514
1358 #define KEY_CLEAR         0515
1359 #define KEY_EOS           0516
1360 #define KEY_EOL           0517
1361 #define KEY_SF            0520
1362 #define KEY_SR            0521
1363 #define KEY_NPAGE         0522
1364 #define KEY_PPAGE         0523
1365 #define KEY_STAB          0524
1366 #define KEY_CTAB          0525
1367 #define KEY_CATAB         0526
1368 #define KEY_ENTER         0527
1369 #define KEY_SRESET        0530
1370 #define KEY_RESET          0531
1371 #define KEY_PRINT          0532
1372 #define KEY_LL             0533
1373 #define KEY_A1             0534
1374 #define KEY_A3             0535
1375 #define KEY_B2             0536
1376 #define KEY_C1             0537
1377 #define KEY_C3             0540
1378 #define KEY_BTAB           0541
1379 #define KEY_BEG            0542
1380 #define KEY_CANCEL         0543
1381 #define KEY_CLOSE          0544

```

```

1382     #define KEY_COMMAND      0545
1383     #define KEY_COPY         0546
1384     #define KEY_CREATE        0547
1385     #define KEY_END          0550
1386     #define KEY_EXIT          0551
1387     #define KEY_FIND          0552
1388     #define KEY_HELP          0553
1389     #define KEY_MARK          0554
1390     #define KEY_MESSAGE        0555
1391     #define KEY_MOVE          0556
1392     #define KEY_NEXT          0557
1393     #define KEY_OPEN          0560
1394     #define KEY_OPTIONS        0561
1395     #define KEY_PREVIOUS       0562
1396     #define KEY_REDO           0563
1397     #define KEY_REFERENCE      0564
1398     #define KEY_REFRESH        0565
1399     #define KEY_REPLACE        0566
1400     #define KEY_RESTART        0567
1401     #define KEY_RESUME         0570
1402     #define KEY_SAVE           0571
1403     #define KEY_SBEG           0572
1404     #define KEY_SCANCEL         0573
1405     #define KEY_SCOMMAND        0574
1406     #define KEY_SCOPY           0575
1407     #define KEY_SCREATE         0576
1408     #define KEY_SDC            0577
1409     #define KEY SDL             0600
1410     #define KEY_SELECT         0601
1411     #define KEY_SEND           0602
1412     #define KEY_SEOL            0603
1413     #define KEY_SEXIT          0604
1414     #define KEY_SFIND          0605
1415     #define KEY_SHELP           0606
1416     #define KEY_SHOME          0607
1417     #define KEY_SIC            0610
1418     #define KEY_SLEFT           0611
1419     #define KEY_SMESSAGE        0612
1420     #define KEY_SMOVE           0613
1421     #define KEY_SNEXT           0614
1422     #define KEY_SOPTIONS        0615
1423     #define KEY_SPREVIOUS       0616
1424     #define KEY_SPRINT          0617
1425     #define KEY_SREDO           0620
1426     #define KEY_SREPLACE         0621
1427     #define KEY_SRIGHT          0622
1428     #define KEY_SRsume          0623
1429     #define KEY_SSAVE           0624
1430     #define KEY_SSUSPEND        0625
1431     #define KEY_SUNDO           0626
1432     #define KEY_SUSPEND         0627
1433     #define KEY_UNDO            0630
1434     #define KEY_MOUSE           0631
1435     #define KEY_RESIZE          0632
1436     #define KEY_MAX            0777
1437
1438     #define PAIR_NUMBER(a)    (((a)&A_COLOR)>>8)
1439     #define NCURSES_BITS(mask,shift)      ((mask)<<((shift)+8))
1440     #define A_CHARTEXT          (NCURSES_BITS(1UL,0)-1UL)
1441     #define A_NORMAL            0L
1442     #define NCURSES_ATTR_SHIFT   8
1443     #define A_COLOR             NCURSES_BITS(((1UL)<<8)-1UL,0)
1444     #define A_BLINK              NCURSES_BITS(1UL,11)
1445     #define A_DIM                NCURSES_BITS(1UL,12)

```

```

1446 #define A_BOLD    NCURSES_BITS(1UL,13)
1447 #define A_ALTCHARSET    NCURSES_BITS(1UL,14)
1448 #define A_INVIS   NCURSES_BITS(1UL,15)
1449 #define A_PROTECT    NCURSES_BITS(1UL,16)
1450 #define A_HORIZONTAL    NCURSES_BITS(1UL,17)
1451 #define A_LEFT    NCURSES_BITS(1UL,18)
1452 #define A_LOW     NCURSES_BITS(1UL,19)
1453 #define A_RIGHT   NCURSES_BITS(1UL,20)
1454 #define A_TOP     NCURSES_BITS(1UL,21)
1455 #define A_VERTICAL    NCURSES_BITS(1UL,22)
1456 #define A_STANDOUT    NCURSES_BITS(1UL,8)
1457 #define A_UNDERLINE    NCURSES_BITS(1UL,9)
1458 #define COLOR_PAIR(n)  NCURSES_BITS(n,0)
1459 #define A_ATTRIBUTES    NCURSES_BITS(~(1UL-1UL),0)
1460
1461 extern int addch(const chtype);
1462 extern int addchnstr(const chtype *, int);
1463 extern int addchstr(const chtype *);
1464 extern int addnstr(const char *, int);
1465 extern int addstr(const char *);
1466 extern int attroff(int);
1467 extern int attron(int);
1468 extern int attrset(int);
1469 extern int attr_get(attr_t *, short *, void *);
1470 extern int attr_off(attr_t, void *);
1471 extern int attr_on(attr_t, void *);
1472 extern int attr_set(attr_t, short, void *);
1473 extern int baudrate(void);
1474 extern int beep(void);
1475 extern int bkgd(chtype);
1476 extern void bkgdset(chtype);
1477 extern int border(chtype, chtype, chtype, chtype, chtype,
1478                   chtype,
1479                   chtype);
1480 extern int box(WINDOW *, chtype, chtype);
1481 extern bool can_change_color(void);
1482 extern int cbreak(void);
1483 extern int chgat(int, attr_t, short, const void *);
1484 extern int clear(void);
1485 extern int clearok(WINDOW *, bool);
1486 extern int clrbot(void);
1487 extern int clrtoeo(void);
1488 extern int color_content(short, short *, short *, short *);
1489 extern int color_set(short, void *);
1490 extern int copywin(const WINDOW *, WINDOW *, int, int, int, int,
1491                   int,
1492                   int);
1493 extern int curs_set(int);
1494 extern int def_prog_mode(void);
1495 extern int def_shell_mode(void);
1496 extern int delay_output(int);
1497 extern int delch(void);
1498 extern void delscreen(SCREEN *);
1499 extern int delwin(WINDOW *);
1500 extern int deleteln(void);
1501 extern WINDOW *derwin(WINDOW *, int, int, int, int);
1502 extern int doupdate(void);
1503 extern WINDOW *dupwin(WINDOW *);
1504 extern int echo(void);
1505 extern int echochar(const chtype);
1506 extern int erase(void);
1507 extern int endwin(void);
1508 extern char erasechar(void);
1509 extern void filter(void);

```

```

1510     extern int flash(void);
1511     extern int flushinp(void);
1512     extern chtype getbkgd(WINDOW *);
1513     extern int getch(void);
1514     extern int getnstr(char *, int);
1515     extern int getstr(char *);
1516     extern WINDOW *getwin(FILE *);
1517     extern int halfdelay(int);
1518     extern bool has_colors(void);
1519     extern bool has_ic(void);
1520     extern bool has_il(void);
1521     extern int hline(chtype, int);
1522     extern void idcok(WINDOW *, bool);
1523     extern int idlok(WINDOW *, bool);
1524     extern void immedok(WINDOW *, bool);
1525     extern chtype inch(void);
1526     extern int inchnstr(chtype *, int);
1527     extern int inchstr(chtype *);
1528     extern WINDOW *initscr(void);
1529     extern int init_color(short, short, short, short);
1530     extern int init_pair(short, short, short);
1531     extern int innstr(char *, int);
1532     extern int insch(chtype);
1533     extern int insdelln(int);
1534     extern int insertln(void);
1535     extern int insnstr(const char *, int);
1536     extern int insstr(const char *);
1537     extern int instr(char *);
1538     extern int intrflush(WINDOW *, bool);
1539     extern bool isendwin(void);
1540     extern bool is_linetouched(WINDOW *, int);
1541     extern bool is_wintouched(WINDOW *);
1542     extern const char *keyname(int);
1543     extern int keypad(WINDOW *, bool);
1544     extern char killchar(void);
1545     extern int leaveok(WINDOW *, bool);
1546     extern char *longname(void);
1547     extern int meta(WINDOW *, bool);
1548     extern int move(int, int);
1549     extern int mvaddch(int, int, const chtype);
1550     extern int mvaddchnstr(int, int, const chtype *, int);
1551     extern int mvaddchstr(int, int, const chtype *);
1552     extern int mvaddnstr(int, int, const char *, int);
1553     extern int mvaddstr(int, int, const char *);
1554     extern int mvchgat(int, int, int, attr_t, short, const void *);
1555     extern int mvcur(int, int, int, int);
1556     extern int mvdelch(int, int);
1557     extern int mvderwin(WINDOW *, int, int);
1558     extern int mvgetch(int, int);
1559     extern int mvgetnstr(int, int, char *, int);
1560     extern int mvgetstr(int, int, char *);
1561     extern int mvhline(int, int, chtype, int);
1562     extern chtype mvinch(int, int);
1563     extern int mvinchnstr(int, int, chtype *, int);
1564     extern int mvinchstr(int, int, chtype *);
1565     extern int mvinnstr(int, int, char *, int);
1566     extern int mvinsch(int, int, chtype);
1567     extern int mvinsnstr(int, int, const char *, int);
1568     extern int mvinsnstr(int, int, const char *);
1569     extern int mvinstr(int, int, char *);
1570     extern int mvprintw(int, int, char *, ...);
1571     extern int mvscanw(int, int, const char *, ...);
1572     extern int mvvline(int, int, chtype, int);
1573     extern int mvwaddch(WINDOW *, int, int, const chtype);

```

```

1574 extern int mvwaddchnstr(WINDOW *, int, int, const chtype *, int);
1575 extern int mvwaddchstr(WINDOW *, int, int, const chtype *);
1576 extern int mvwaddnstr(WINDOW *, int, int, const char *, int);
1577 extern int mvwaddstr(WINDOW *, int, int, const char *);
1578 extern int mvwchgat(WINDOW *, int, int, int, attr_t, short, const void
1579 * );
1580 extern int mvwdelch(WINDOW *, int, int);
1581 extern int mvwgetch(WINDOW *, int, int);
1582 extern int mvwgetnstr(WINDOW *, int, int, char *, int);
1583 extern int mvwgetstr(WINDOW *, int, int, char *);
1584 extern int mvwhline(WINDOW *, int, int, chtype, int);
1585 extern int mvwin(WINDOW *, int, int);
1586 extern chtype mvwinch(WINDOW *, int, int);
1587 extern int mvwinchnstr(WINDOW *, int, int, chtype *, int);
1588 extern int mvwinchstr(WINDOW *, int, int, chtype *);
1589 extern int mvwinnstr(WINDOW *, int, int, char *, int);
1590 extern int mvwinsch(WINDOW *, int, int, chtype);
1591 extern int mvwinsnstr(WINDOW *, int, int, const char *, int);
1592 extern int mvwinsstr(WINDOW *, int, int, const char *);
1593 extern int mvwinstr(WINDOW *, int, int, char *);
1594 extern int mvwprintw(WINDOW *, int, int, char *, ...);
1595 extern int mvwscanw(WINDOW *, int, int, const char *, ...);
1596 extern int mvwvline(WINDOW *, int, int, chtype, int);
1597 extern int napms(int);
1598 extern WINDOW *newpad(int, int);
1599 extern SCREEN *newterm(const char *, FILE *, FILE *);
1600 extern WINDOW *newwin(int, int, int, int);
1601 extern int nl(void);
1602 extern int nocbreak(void);
1603 extern int nodelay(WINDOW *, bool);
1604 extern int noecho(void);
1605 extern int nonl(void);
1606 extern void noqiflush(void);
1607 extern int noraw(void);
1608 extern int notimeout(WINDOW *, bool);
1609 extern int overlay(const WINDOW *, WINDOW *);
1610 extern int overwrite(const WINDOW *, WINDOW *);
1611 extern int pair_content(short, short *, short *);
1612 extern int pechochar(WINDOW *, chtype);
1613 extern int pnoutrefresh(WINDOW *, int, int, int, int, int, int);
1614 extern int prefresh(WINDOW *, int, int, int, int, int, int);
1615 extern int printw(char *, ...);
1616 extern int putwin(WINDOW *, FILE *);
1617 extern void qiflush(void);
1618 extern int raw(void);
1619 extern int redrawwin(WINDOW *);
1620 extern int refresh(void);
1621 extern int resetty(void);
1622 extern int reset_prog_mode(void);
1623 extern int reset_shell_mode(void);
1624 extern int ripoffline(int, int (*init) (WINDOW *, int)
1625 );
1626 extern int savetty(void);
1627 extern int scanw(const char *, ...);
1628 extern int scr_dump(const char *);
1629 extern int scr_init(const char *);
1630 extern int scrl(int);
1631 extern int scroll(WINDOW *);
1632 extern int scrolllok(WINDOW *, typedef unsigned char bool);
1633 extern int scr_restore(const char *);
1634 extern int scr_set(const char *);
1635 extern int setsqrreg(int, int);
1636 extern SCREEN *set_term(SCREEN *);
1637 extern int slk_attroff(const typedef unsigned long int chtype);
```

```

1638     extern int slk_attron(const typedef unsigned long int chtype);
1639     extern int slk_attrset(const typedef unsigned long int chtype);
1640     extern int slk_attr_set(const typedef chtype attr_t, short, void *);
1641     extern int slk_clear(void);
1642     extern int slk_color(short);
1643     extern int slk_init(int);
1644     extern char *slk_label(int);
1645     extern int slk_noutrefresh(void);
1646     extern int slk_refresh(void);
1647     extern int slk_restore(void);
1648     extern int slk_set(int, const char *, int);
1649     extern int slk_touch(void);
1650     extern int standout(void);
1651     extern int standend(void);
1652     extern int start_color(void);
1653     extern WINDOW *subpad(WINDOW *, int, int, int, int);
1654     extern WINDOW *subwin(WINDOW *, int, int, int, int);
1655     extern int syncok(WINDOW *, typedef unsigned char bool);
1656     extern typedef unsigned long int chtype termattrs(void);
1657     extern char *termname(void);
1658     extern void timeout(int);
1659     extern int typeahead(int);
1660     extern int ungetch(int);
1661     extern int untouchwin(WINDOW *);
1662     extern void use_env(typedef unsigned char bool);
1663     extern int vidattr(typedef unsigned long int chtype);
1664     extern int vidputs(typedef unsigned long int chtype,
1665                         int (*vidputs_int) (int)
1666                         );
1667     extern int vline(typedef unsigned long int chtype, int);
1668     extern int vwprintw(WINDOW *, char *, typedef void *va_list);
1669     extern int vw_printw(WINDOW *, const char *, typedef void *va_list);
1670     extern int vwscanw(WINDOW *, const char *, typedef void *va_list);
1671     extern int vw_scanw(WINDOW *, const char *, typedef void *va_list);
1672     extern int waddch(WINDOW *, const typedef unsigned long int chtype);
1673     extern int waddchnstr(WINDOW *, const typedef unsigned long int chtype
1674                           *,
1675                           int);
1676     extern int waddchstr(WINDOW *, const typedef unsigned long int chtype
1677                           *);
1678     extern int waddnstr(WINDOW *, const char *, int);
1679     extern int waddstr(WINDOW *, const char *);
1680     extern int wattroff(WINDOW *, int);
1681     extern int wattroff(WINDOW *, int);
1682     extern int wattrset(WINDOW *, int);
1683     extern int wattr_get(WINDOW *, attr_t *, short *, void *);
1684     extern int wattr_on(WINDOW *, typedef chtype attr_t, void *);
1685     extern int wattr_off(WINDOW *, typedef chtype attr_t, void *);
1686     extern int wattr_set(WINDOW *, typedef chtype attr_t, short, void *);
1687     extern int wbkgd(WINDOW *, typedef unsigned long int chtype);
1688     extern void wbkgdset(WINDOW *, typedef unsigned long int chtype);
1689     extern int wborder(WINDOW *, typedef unsigned long int chtype,
1690                         typedef unsigned long int chtype,
1691                         typedef unsigned long int chtype,
1692                         typedef unsigned long int chtype,
1693                         typedef unsigned long int chtype,
1694                         typedef unsigned long int chtype,
1695                         typedef unsigned long int chtype,
1696                         typedef unsigned long int chtype);
1697     extern int wchgat(WINDOW *, int, typedef chtype attr_t, short,
1698                       const void *);
1699     extern int wclear(WINDOW *);
1700     extern int wclrtobot(WINDOW *);
1701     extern int wclrtoeol(WINDOW *);

```

```

1702 extern int wcolor_set(WINDOW *, short, void *);
1703 extern void wcursyncup(WINDOW *);
1704 extern int wdelch(WINDOW *);
1705 extern int wdeleteln(WINDOW *);
1706 extern int wechochar(WINDOW *, const typedef unsigned long int chtype);
1707 extern int werase(WINDOW *);
1708 extern int wgetch(WINDOW *);
1709 extern int wgetnstr(WINDOW *, char *, int);
1710 extern int wgetstr(WINDOW *, char *);
1711 extern int whline(WINDOW *, typedef unsigned long int chtype, int);
1712 extern typedef unsigned long int chtype winch(WINDOW *);
1713 extern int winchnstr(WINDOW *, chtype *, int);
1714 extern int winchstr(WINDOW *, chtype *);
1715 extern int winnstr(WINDOW *, char *, int);
1716 extern int winsch(WINDOW *, typedef unsigned long int chtype);
1717 extern int winsdelln(WINDOW *, int);
1718 extern int wininsertln(WINDOW *);
1719 extern int winsnstr(WINDOW *, const char *, int);
1720 extern int winsstr(WINDOW *, const char *);
1721 extern int winstr(WINDOW *, char *);
1722 extern int wmove(WINDOW *, int, int);
1723 extern int wnoutrefresh(WINDOW *);
1724 extern int wprintf(WINDOW *, char *, ...);
1725 extern int wredrawln(WINDOW *, int, int);
1726 extern int wrefresh(WINDOW *);
1727 extern int wscanw(WINDOW *, const char *, ...);
1728 extern int wscrell(WINDOW *, int);
1729 extern int wsetscreg(WINDOW *, int, int);
1730 extern int wstandout(WINDOW *);
1731 extern int wstandend(WINDOW *);
1732 extern void wsyncdown(WINDOW *);
1733 extern void wsyncup(WINDOW *);
1734 extern void wtimeout(WINDOW *, int);
1735 extern int wtouchln(WINDOW *, int, int, int);
1736 extern int wvline(WINDOW *, typedef unsigned long int chtype, int);
1737 extern char *unctrl(typedef unsigned long int chtype);
1738 extern int COLORS(void);
1739 extern int COLOR_PAIRS(void);
1740 extern chtype acs_map(void);
1741 extern WINDOW *curscr(void);
1742 extern WINDOW *stdscr(void);
1743 extern int COLS(void);
1744 extern int LINES(void);
1745 extern int touchline(WINDOW *, int, int);
1746 extern int touchwin(WINDOW *);

```

## 14.6.2 term.h

```

1747 extern int putp(const char *);
1748 extern int tigetflag(const char *);
1749 extern int tigetnum(const char *);
1750 extern char *tigetstr(const char *);
1751 extern char *tparm(const char *, ...);
1752 extern TERMINAL *set_curterm(TERMINAL *);
1753 extern int del_curterm(TERMINAL *);
1754 extern int restartterm(char *, int, int *);
1755 extern int setupterm(char *, int, int *);
1756 extern char *tgetstr(char *, char **);
1757 extern char *tgoto(const char *, int, int);
1758 extern int tgetent(char *, const char *);
1759 extern int tgetflag(char *);
1760 extern int tgetnum(char *);
1761 extern int tputs(const char *, int, int (*putcproc) (int))

```

```
1763     );
1764     extern TERMINAL *cur_term(void);
```

## 14.7 Interfaces for libutil

1765 Table 14-6 defines the library name and shared object name for the libutil library

1766 **Table 14-6 libutil Definition**

1767	Library:	libutil
	SONAME:	libutil.so.1

1768 The behavior of the interfaces in this library is specified by the following specifications:  
1769

1770 [LSB] This Specification

### 14.7.1 Utility Functions

#### 14.7.1.1 Interfaces for Utility Functions

1772 An LSB conforming implementation shall provide the generic functions for Utility  
1773 Functions specified in Table 14-7, with the full mandatory functionality as described  
1774 in the referenced underlying specification.

1775 **Table 14-7 libutil - Utility Functions Function Interfaces**

1776	forkpty [LSB]	login [LSB]	login_tty [LSB]	logout [LSB]
	logwtmp [LSB]	openpty [LSB]		

## 14.8 Interface Definitions for libutil

1777 The interfaces defined on the following pages are included in libutil and are defined  
1778 by this specification. Unless otherwise noted, these interfaces shall be included in the  
1779 source standard.

1780 Other interfaces listed in Section 14.7 shall behave as described in the referenced  
1781 base document.

## forkpty

### Name

1782      `forkpty` — Create a new process attached to an available pseudo-terminal

### Synopsis

```
1783 #include <pty.h>
1784 int forkpty(int * amaster, char * name, struct termios * termpt, struct winsize
1785 * wintp);
```

### Description

1786 The `forkpty()` function shall find and open a pseudo-terminal device pair in the  
 1787 same manner as the `openpty()` function. If a pseudo-terminal is available,  
 1788 `forkpty()` shall create a new process in the same manner as the `fork()` function,  
 1789 and prepares the new process for login in the same manner as `login_tty()`.

1790 If `termpt` is not null, it shall refer to a `termios` structure that shall be used to initialize  
 1791 the characteristics of the slave device. If `wintp` is not null, it shall refer to a `winsize`  
 1792 structure used to initialize the window size of the slave device.

### Return Value

1793 On success, the parent process shall return the process id of the child, and the child  
 1794 shall return 0. On error, no new process shall be created, -1 shall be returned, and  
 1795 `errno` shall be set appropriately. On success, the parent process shall receive the file  
 1796 descriptor of the master side of the pseudo-terminal in the location referenced by  
 1797 `amaster`, and, if `name` is not NULL, the filename of the slave device in `name`.

### Errors

1798 `EAGAIN`

1799      Unable to create a new process.

1800 `ENOENT`

1801      There are no available pseudo-terminals.

1802 `ENOMEM`

1803      Insufficient memory was available.

## login

### Name

1804            login — login utility function

### Synopsis

1805            #include <utmp.h>  
 1806            void login (struct utmp \* *ut* );

### Description

1807            The *login()* function shall update the user accounting databases. The *ut* parameter  
 1808            shall reference a *utmp* structure for all fields except the following:

- 1809            1. The *ut\_type* field shall be set to *USER\_PROCESS*.
- 1810            2. The *ut\_pid* field shall be set to the process identifier for the current process.
- 1811            3. The *ut\_line* field shall be set to the name of the controlling terminal device.  
                   The name shall be found by examining the device associated with the standard  
                   input, output and error streams in sequence, until one associated with a  
                   terminal device is found. If none of these streams refers to a terminal device,  
                   the *ut\_line* field shall be set to "???". If the terminal device is in the */dev*  
                   directory hierarchy, the *ut\_line* field shall not contain the leading *"/dev/"*,  
                   otherwise it shall be set to the final component of the pathname of the device. If  
                   the user accounting database imposes a limit on the size of the *ut\_line* field, it  
                   shall truncate the name, but any such limit shall not be smaller than  
                   *UT\_LINESIZE* (including a terminating null character).

### Return Value

1821            None

### Errors

1822            None

## **login\_tty**

### **Name**

1823      `login_tty` — Prepare a terminal for login

### **Synopsis**

1824      `#include <utmp.h>`  
 1825      `int login_tty ( int fdr );`

### **Description**

1826      The `login_tty()` function shall prepare the terminal device referenced by the file  
 1827      descriptor `fdr`. This function shall create a new session, make the terminal the  
 1828      controlling terminal for the current process, and set the standard input, output, and  
 1829      error streams of the current process to the terminal. If `fdr` is not the standard input,  
 1830      output or error stream, then `login_tty()` shall close `fdr`.

### **Return Value**

1831      On success, `login_tty()` shall return zero; otherwise -1 is returned, and `errno` shall  
 1832      be set appropriately.

### **Errors**

1833      ENOTTY

1834      `fdr` does not refer to a terminal device.

## **logout**

### **Name**

1835      `logout` — logout utility function

### **Synopsis**

1836      `#include <utmp.h>`  
 1837      `int logout ( const char * line );`

### **Description**

1838      Given the device `line`, the `logout()` function shall search the user accounting  
 1839      database which is read by `getutent()` for an entry with the corresponding line, and  
 1840      with the type of `USER_PROCESS`. If a corresponding entry is located, it shall be  
 1841      updated as follows:

- 1842      1. The `ut_name` field shall be set to zeroes (`UT_NAMESIZE` NUL bytes).
- 1843      2. The `ut_host` field shall be set to zeroes (`UT_HOSTSIZE` NUL bytes).
- 1844      3. The `ut_tv` shall be set to the current time of day.
- 1845      4. The `ut_type` field shall be set to `DEAD_PROCESS`.

### **Return Value**

1846      On success, the `logout()` function shall return non-zero. Zero is returned if there  
 1847      was no entry to remove, or if the `utmp` file could not be opened or updated.

## logwtmp

### Name

1848 logwtmp — append an entry to the wtmp file

### Synopsis

```
1849 #include <utmp.h>
1850 void logwtmp (const char * line , const char * name , const char * host );
```

### Description

1851 If the process has permission to update the user accounting databases, the  
 1852 `logwtmp()` function shall append a record to the user accounting database that  
 1853 records all logins and logouts. The record to be appended shall be constructed as  
 1854 follows:

- 1855 1. The `ut_line` field shall be initialized from `line`. If the user accounting  
 1856 database imposes a limit on the size of the `ut_line` field, it shall truncate the  
 1857 value, but any such limit shall not be smaller than `UT_LINESIZE` (including a  
 1858 terminating null character).
- 1859 2. The `ut_name` field shall be initialized from `name`. If the user accounting  
 1860 database imposes a limit on the size of the `ut_name` field, it shall truncate the  
 1861 value, but any such limit shall not be smaller than `UT_NAMESIZE` (including a  
 1862 terminating null character).
- 1863 3. The `ut_host` field shall be initialized from `host`. If the user accounting  
 1864 database imposes a limit on the size of the `ut_host` field, it shall truncate the  
 1865 value, but any such limit shall not be smaller than `UT_HOSTSIZE` (including a  
 1866 terminating null character).
- 1867 4. If the `name` parameter does not refer to an empty string (i.e. " "), the `ut_type`  
 1868 field shall be set to `USER_PROCESS`; otherwise the `ut_type` field shall be set to  
 1869 `DEAD_PROCESS`.
- 1870 5. The `ut_id` field shall be set to the process identifier for the current process.
- 1871 6. The `ut_tv` field shall be set to the current time of day.

1872 **Note:** If a process does not have write access to the the user accounting database, the  
 1873 `logwtmp()` function will not update it. Since the function does not return any value, an  
 1874 application has no way of knowing whether it succeeded or failed.

### Return Value

1875 None.

## openpty

### Name

1876 openpty — find and open an available pseudo-terminal

### Synopsis

```
1877 #include <pty.h>
1878 int openpty(int *amaster, int *aslave, char *name, struct termios *termp,
1879 struct winsize *winp);
```

### Description

1880 The `openpty()` function shall find an available pseudo-terminal and return file  
 1881 descriptors for the master and slave devices in the locations referenced by `amaster`  
 1882 and `aslave` respectively. If `name` is not NULL, the filename of the slave shall be  
 1883 placed in the user supplied buffer referenced by `name`. If `termp` is not NULL, it shall  
 1884 point to a `termios` structure used to initialize the terminal parameters of the slave  
 1885 pseudo-terminal device. If `winp` is not NULL, it shall point to a `winsize` structure  
 1886 used to initialize the window size parameters of the slave pseudo-terminal device.

### Return Value

1887 On success, zero is returned. On error, -1 is returned, and `errno` is set appropriately.

### Errors

1888 ENOENT

1889 There are no available pseudo-terminals.

## V Commands and Utilities

# 15 Commands and Utilities

## 15.1 Commands and Utilities

1 An LSB conforming implementation shall provide the commands and utilities as  
2 described in Table 15-1, with at least the behavior described as mandatory in the  
3 referenced underlying specification, with the following exceptions:

- 4 1. If any operand (except one which follows --) starts with a hyphen, the  
5 behavior is unspecified.

6 **Rationale (Informative):** Applications should place options before operands, or use --,  
7 as needed. This text is needed because, by default, GNU option parsing differs  
8 from POSIX, unless the environment variable POSIXLY\_CORRECT is set. For  
9 example, **ls . -a** in GNU **ls** means to list the current directory, showing all files (that  
10 is, " ." is an operand and -a is an option). In POSIX, " ." and -a are both operands,  
11 and the command means to list the current directory, and also the file named -a.  
12 Suggesting that applications rely on the setting of the POSIXLY\_CORRECT  
13 environment variable, or try to set it, seems worse than just asking the applications  
14 to invoke commands in ways which work with either the POSIX or GNU  
15 behaviors.

16 **Table 15-1 Commands And Utilities**

[ [1]	dmesg [2]	id [1]	mount [2]	sort [1]
ar [2]	du [2]	install [2]	msgfmt [2]	split [1]
at [2]	echo [2]	install_initd [2]	mv [1]	strip [1]
awk [2]	ed [1]	ipcrm [2]	newgrp [2]	stty [1]
basename [1]	egrep [2]	ipcs [2]	nice [1]	su [2]
batch [2]	env [1]	join [1]	nl [1]	sync [2]
bc [2]	expand [1]	kill [1]	nohup [1]	tail [1]
cat [1]	expr [1]	killall [2]	od [2]	tar [2]
chfn [2]	false [1]	ln [1]	passwd [2]	tee [1]
chgrp [1]	fgrep [2]	locale [1]	paste [1]	test [1]
chmod [1]	file [2]	localedef [1]	patch [2]	time [1]
chown [1]	find [2]	logger [1]	pathchk [1]	touch [1]
chsh [2]	fold [1]	logname [1]	pax [1]	tr [1]
cksum [1]	fuser [2]	lp [1]	pidof [2]	true [1]
cmp [1]	gencat [1]	lpr [2]	pr [1]	tsort [1]
col [2]	getconf [1]	ls [2]	printf [1]	tty [1]
comm [1]	gettext [2]	lsb_release [2]	ps [1]	umount [2]
cp [1]	grep [2]	m4 [2]	pwd [1]	uname [1]
cpio [2]	groupadd [2]	mailx [1]	remove_initd	unexpand [1]

			[2]	
17	crontab [2]	groupdel [2]	make [1]	renice [2]
	csplit [1]	groupmod [2]	man [1]	rm [1]
	cut [2]	groups [2]	md5sum [2]	rmdir [1]
	date [1]	gunzip [2]	mkdir [1]	sed [2]
	dd [1]	gzip [2]	mkfifo [1]	sendmail [2]
	df [2]	head [1]	mknod [2]	sh [2]
	diff [1]	hostname [2]	mktemp [2]	shutdown [2]
	dirname [1]	iconv [1]	more [2]	sleep [1]

18        *Referenced Specification(s)*

19        [1]. ISO POSIX (2003)

20        [2]. This Specification

21        An LSB conforming implementation shall provide the shell built in utilities as  
22        described in Table 15-2, with at least the behavior described as mandatory in the  
23        referenced underlying specification, with the following exceptions:

- 24        1. The built in commands and utilities shall be provided by the
- sh**
- utility itself,
- 
- 25            and need not be implemented in a manner so that they can be accessed via the
- 
- 26            exec family of functions as defined in ISO POSIX (2003) and should not be
- 
- 27            invoked directly by those standard utilities that execute other utilities (
- env**
- ,
- 
- 28
- find**
- ,
- nice**
- ,
- nohup**
- ,
- time**
- ,
- xargs**
- ).

29        **Rationale (Informative):** Since the built in utilities must affect the environment of the  
30            calling process, they have no effect when executed as a file.31        **Table 15-2 Built In Utilities**

32        cd [1]	getopts [1]	read [1]	umask [1]	wait [1]
------------------	-------------	----------	-----------	----------

33        *Referenced Specification(s)*

34        [1]. ISO POSIX (2003)

## 15.2 Command Behavior

35        This section contains descriptions for commands and utilities whose specified  
36            behavior in the LSB contradicts or extends the standards referenced. It also contains  
37            commands and utilities only required by the LSB and not specified by other  
38            standards.

**ar****Name**

39        ar – create and maintain library archives (DEPRECATED)

**Description**

40        ar is deprecated from the LSB and is expected to disappear from a future version of  
41        the LSB.

42        **Rationale:** The LSB generally does not include software development utilities nor does it  
43        specify .o and .a file formats.

44        ar is as specified in ISO POSIX (2003) but with differences as listed below.

**Differences**

45        -T

46        -C

47              need not be accepted.

48        -l

49              has unspecified behavior.

50        -q

51              has unspecified behavior; using -r is suggested.

## at

### Name

52           at — examine or delete jobs for later execution

### Description

53           at is as specified in ISO POSIX (2003) but with differences as listed below.

### Differences

54           Options

55           -d

56                 is functionally equivalent to the -r option specified in ISO POSIX (2003).

57           -r

58                 need not be supported, but the '-d' option is equivalent.

59           -t time

60                 need not be supported.

61           Optional Control Files

62           The implementation shall support the XSI optional behavior for access control;  
63           however the files at.allow and at.deny may reside in /etc rather than  
64           /usr/lib/cron.

## awk

### Name

65           awk — pattern scanning and processing language

### Description

66           awk is as specified in ISO POSIX (2003) but with differences as listed below.

### Differences

67           Certain aspects of internationalized regular expressions are optional; see  
68           Internationalization and Regular Expressions.

**batch****Name**

69      **batch** — schedule commands to be executed in a batch queue

**Description**

70      The specification for **batch** is as specified in ISO POSIX (2003), but with differences  
71      as listed below.

**Optional Control Files**

73      The implementation shall support the XSI optional behavior for access control;  
74      however the files `at.allow` and `at.deny` may reside in `/etc` rather than  
75      `/usr/lib/cron`.

**bc****Name**

76      **bc** — an arbitrary precision calculator language

**Description**

77      **bc** is as specified in ISO POSIX (2003) but with extensions as listed below.

**Extensions**

78      The bc language may be extended in an implementation defined manner. If an  
79      implementation supports extensions, it shall also support the additional options:

80      `-s | --standard`  
81            processes exactly the POSIX **bc** language.  
82      `-w | --warn`  
83            gives warnings for extensions to POSIX bc.

**chfn****Name**

84      **chfn** — change user name and information

**Synopsis**

85      **chfn** [ `-f full_name` ] [ `-h home_phone` ] [ `user` ]

**Description**

86      **chfn** shall update the user database. An unprivileged user may only change the  
87      fields for their own account, a user with appropriate privileges may change the  
88      fields for any account.

89      The fields `full_name` and `home_phone` may contain any character except:

any control character  
comma  
colon  
equal sign

If none of the options are selected, **chfn** operates in an interactive fashion. The prompts and expected input in interactive mode are unspecified and should not be relied upon.

As it is possible for the system to be configured to restrict which fields a non-privileged user is permitted to change, applications should be written to gracefully handle these situations.

## Standard Options

-f *full\_name*  
sets the user's full name.

-h *home\_phone*  
sets the user's home phone number.

## Future Directions

The following two options are expected to be added in a future version of the LSB:

-o office  
sets the user's office room number.

-p office\_phone  
sets the user's office phone number.

Note that some implementations contain a "-o other" option which specifies an additional field called "other". Traditionally, this field is not subject to the constraints about legitimate characters in fields. Also, one traditionally shall have appropriate privileges to change the other field. At this point there is no consensus about whether it is desirable to specify the other field; applications may wish to avoid using it.

The "-w work\_phone" field found in some implementations should be replaced by the "-p office\_phone" field. The "-r room\_number" field found in some implementations is the equivalent of the "-o office" option mentioned above; which one of these two options to specify will depend on implementation experience and the decision regarding the other field.

## chsh

### Name

117 chsh – change login shell

### Synopsis

118 **chsh** [-s *login\_shell*] [*user*]

### Description

119 **chsh** changes the user login shell. This determines the name of the user's initial login  
120 command. An unprivileged user may only change the login shell for their own  
121 account, a user with appropriate privilege may change the login shell for any  
122 account specified by *user*.

123 Unless the user has appropriate privilege, the initial login command name shall be  
124 one of those listed in /etc/shells. The *login\_shell* shall be the absolute path (i.e.  
125 it must start with '/') to an executable file. Accounts which are restricted (in an  
126 implementation-defined manner) may not change their login shell.

127 If the -s option is not selected, **chsh** operates in an interactive mode. The prompts  
128 and expected input in this mode are unspecified.

### Standard Options

129 -s *login\_shell*

130 sets the login shell.

## col

### Name

131 **col** – filter reverse line feeds from input

### Description

132 **col** is as specified in SUSv2 but with differences as listed below.

### Differences

133 The -p option has unspecified behavior.

134 **Note:** Although **col** is shown as legacy in SUSv2, it is not (yet) deprecated in the LSB.

## cpio

### Name

135 cpio — copy file archives in and out

### Description

136 cpio is as specified in ISO POSIX (2003), but with differences as listed below.

### Differences

137 Some elements of the Pattern Matching Notation are optional; see  
138 Internationalization and Pattern Matching Notation.

## crontab

### Name

139 crontab — maintain crontab files for individual users

### Synopsis

140 `crontab [-u user] file crontab [-u user] {-l | -r | -e}`

### Description

141 crontab is as specified in ISO POSIX (2003), but with differences as listed below.

### Optional Control Files

142 The implementation shall support the XSI optional behavior for access control;  
143 however the files `cron.allow` and `cron.deny` may reside in `/etc` rather than  
144 `/usr/lib/cron`.

## cut

### Name

145 cut — split a file into sections determined by context lines

### Description

146 cut is as specified in ISO POSIX (2003), but with differences as listed below.

### Differences

147 `-n`  
148 has unspecified behavior.

**df****Name**

149      df — report file system disk space usage

**Description**

150      The **df** command shall behave as specified in ISO POSIX (2003), but with differences  
151      as listed below.

**Differences****Options**

153      If the **-k** option is not specified, disk space is shown in unspecified units. If the **-P**  
154      option is specified, the size of the unit shall be printed on the header line in the  
155      format "%4s-blocks". Applications should specify **-k**.

156      The XSI option **-t** has unspecified behavior. Applications should not specify **-t**.

157      **Rationale:** The most common implementation of **df** uses the **-t** option for a different  
158      purpose (restricting output to a particular file system type), and use of **-t** is therefore  
159      non-portable.

**Operand May Identify Special File**

161      If an argument is the absolute file name of a special file containing a mounted file  
162      system, **df** shall show the space available on that file system rather than on the file  
163      system containing the special file (which is typically the root file system).

164      **Note:** In ISO POSIX (2003) the XSI optional behavior permits an operand to name a  
165      special file, but appears to require the operation be performed on the file system  
166      containing the special file. A defect report has been submitted for this case.

## dmesg

### Name

167 **dmesg** — print or control the system message buffer

### Synopsis

168 **dmesg** [-c | -n *level* | -s *bufsize*]

### Description

169 **dmesg** examines or controls the system message buffer. Only a user with  
170 appropriate privileges may modify the system message buffer parameters or  
171 contents.

### Standard Options

172 **-c**

173 If the user has appropriate privilege, clears the system message buffer contents  
174 after printing.

175 **-n *level***

176 If the user has appropriate privilege, sets the level at which logging of messages  
177 is done to the console.

178 **-s *bufsize***

179 uses a buffer of *bufsize* to query the system message buffer. This is 16392 by  
180 default.

## du

### Name

181 **du** — estimate file space usage

### Description

182 **du** is as specified in ISO POSIX (2003), but with differences as listed below.

### Differences

183 If the **-k** option is not specified, disk space is shown in unspecified units.  
184 Applications should specify **-k**.

**echo****Name**

185       **echo** — write arguments to standard output

**Synopsis**

186       **echo** [*string...*]

**Description**

187       The **echo** command is as specified in ISO POSIX (2003), but with the following  
188       differences.

189       Implementations may support implementation-defined options to **echo**. The  
190       behavior of **echo** if any arguments contain backslashes is also implementation  
191       defined.

**Application Usage**

192       Conforming applications should not run **echo** with a first argument starting with a  
193       hyphen, or with any arguments containing backslashes; they should use **printf** in  
194       those cases.

195       **Note:** The behavior specified here is similar to that specified by ISO POSIX (2003)  
196       without the XSI option. However, the LSB strongly recommends conforming  
197       applications not use any options (even if the implementation provides them) while ISO  
198       POSIX (2003) specifies behavior if the first operand is the string *-n*.

**egrep****Name**

199       **egrep** — search a file with an Extended Regular Expression pattern

**Description**

200       **egrep** is equivalent to **grep -E**. For further details, see the specification for **grep**.

**fgrep****Name**

201       **fgrep** — search a file with a fixed pattern

**Description**

202       **fgrep** is equivalent to **grep -F**. For further details, see the specification for **grep**.

## file

### Name

203        **file** — determine file type

### Description

204        **file** is as specified in ISO POSIX (2003), but with differences as listed below.

### Differences

205        The **-M**, **-h**, **-d**, and **-i** options need not be supported.

## find

### Name

206        **find** — search for files in a directory hierarchy

### Description

207        **find** shall behave as specified in ISO POSIX (2003), except as described below.

### Differences

208        **Pattern Matching**

209        Some elements of the Pattern Matching Notation are optional; see  
210        Internationalization and Pattern Matching Notation.

211        **Option and Operand Handling**

212        Options and operands to **find** shall behave as described in ISO POSIX (2003), except  
213        as follows:

214        **-H**

215              need not be supported

216        **-L**

217              need not be supported

218        **-exec** . . . +

219              argument aggregation need not be supported

220        **Rationale:** The **-H** and **-L** options are not yet widely available in implementations of the  
221        **find** command, nor is argument aggregation. A future version of this specification will  
222        require these features be supported.

## fuser

### Name

223 fuser — identify processes using files or sockets

### Description

224 fuser is as specified in ISO POSIX (2003), but with differences as listed below.

### Differences

225 The fuser command is a system administration utility, see Path For System  
226 Administration Utilities.

### Option Differences

228 -c

229 has unspecified behavior.

230 -f

231 has unspecified behavior.

## gettext

### Name

232   **gettext** — retrieve text string from message catalog

### Synopsis

233   **gettext** [options] [*textdomain*] *msgid* **gettext** -s [options] *msgid*...

### Description

234   The **gettext** utility retrieves a translated text string corresponding to string *msgid*  
 235   from a message object generated with **msgfmt** utility.

236   The message object name is derived from the optional argument *textdomain* if  
 237   present, otherwise from the **TEXTDOMAIN** environment variable. If no domain is  
 238   specified, or if a corresponding string cannot be found, **gettext** prints *msgid*.

239   Ordinarily **gettext** looks for its message object in *dirname*/lang/LC\_MESSAGES where  
 240   *dirname* is the implementation-defined default directory and *lang* is the locale  
 241   name. If present, the **TEXTCOMAINDIR** environment variable replaces the *dirname*.

242   This utility interprets C escape sequences such as \t for tab. Use \\ to print a  
 243   backslash. To produce a message on a line of its own, either put a \n at the end of  
 244   *msgid*, or use this command in conjunction with the **printf** utility.

245   When used with the -s option the **gettext** utility behaves like the **echo** utility, except  
 246   that the message corresponding to *msgid* in the selected catalog provides the  
 247   arguments.

### Options

248   -d *domainname*  
 249   --domain=*domainname*  
 250                   PARAMETER translated messages from *domainname*.  
 251   -e  
 252                   Enable expansion of some escape sequences.  
 253   -n  
 254                   Suppress trailing newline.

### Operands

255   The following operands are supported:

256   *textdomain*  
 257                   A domain name used to retrieve the messages.  
 258   *msgid*  
 259                   A key to retrieve the localized message.

### Environment Variables

260   LANGUAGE  
 261                   Specifies one or more locale names.

262            `LANG`  
263                Specifies locale name.  
264            `LC_MESSAGES`  
265                Specifies messaging locale, and if present overrides `LANG` for messages.  
266            `TEXTDOMAIN`  
267                Specifies the text domain name, which is identical to the message object  
268                filename without `.mo` suffix.  
269            `TEXTDOMAINDIR`  
270                Specifies the pathname to the message catalog, and if present replaces the  
271                implementation-defined default directory.

## Exit Status

272            The following exit values are returned:  
273                0  
274                Successful completion.  
275                >0  
276                An error occurred.

# grep

## Name

277            `grep` – print lines matching a pattern

## Description

278            `grep` is as specified in ISO POSIX (2003), but with differences as listed below.

## LSB Differences

279            Certain aspects of regular expression matching are optional; see Internationalization  
280            and Regular Expressions.

## groupadd

### Name

281 groupadd — create a new group

### Synopsis

282 **groupadd** [-g *gid* [-o]] *group*

### Description

283 If the caller has appropriate privilege, the **groupadd** command shall create a new  
284 group named *group*. The group name shall be unique in the group database. If no  
285 *gid* is specified, **groupadd** shall create the new group with a unique group ID.

286 The **groupadd** command is a system administration utility, see Path For System  
287 Administration Utilities.

### Options

288 -g *gid* [-o]

289 The new group shall have group ID *gid*. If the -o option is not used, no other  
290 group shall have this group ID. The value of *gid* shall be non-negative.

## groupdel

### Name

291 groupdel — delete a group

### Synopsis

292 **groupdel** *group*

### Description

293 If the caller has sufficient privilege, the **groupdel** command shall modify the system  
294 group database, deleting the group named *group*. If the group named *group* does  
295 not exist, **groupdel** shall issue a diagnostic message and exit with a non-zero exit  
296 status.

297 The **groupdel** command is a system administration utility, see Path For System  
298 Administration Utilities.

## groupmod

### Name

299        `groupmod` — modify a group

### Synopsis

300        `groupmod [-g gid [-o]] [-n group_name] group`

### Description

301        If the caller has appropriate privilege, the `groupmod` command shall modify the  
302        entry in the system group database corresponding to a group named *group*.

303        The `groupmod` command is a system administration utility, see Path For System  
304        Administration Utilities.

### Options

305        `-g gid [-o]`

306        Modify the group's group ID, setting it to *gid*. If the *-o* option is not used, no  
307        other group shall have this group ID. The value of *gid* shall be non-negative.

308        **Note:** Only the group ID in the database is altered; any files with group ownership set to  
309        the original group ID are unchanged by this modification.

310        `-n group_name`

311        changes the name of the group from *group* to *group\_name*.

## groups

### Name

312        `groups` — display a group

### Synopsis

313        `groups [user]`

### Description

314        The `groups` command shall behave as `id -Gn [user]`, as specified in ISO POSIX  
315        (2003). The optional *user* parameter will display the groups for the named user.

## gunzip

### Name

316        `gunzip` — uncompress files

### Description

317        `gunzip` is equivalent to `gzip -d`. See the specification for `gzip` for further details.

**gzip****Name**

318      **gzip** — compress or expand files

**Synopsis**

319      **gzip** [-cdfhlLnNrtvV19] [-S suffix] [name...]

**Description**

320      The **gzip** command shall attempt to reduce the size of the named files. Whenever  
 321      possible, each file is replaced by one with the extension .gz, while keeping the same  
 322      ownership, modes, access and modification times. If no files are specified, or if a file  
 323      name is -, the standard input is compressed to the standard output. **gzip** shall only  
 324      attempt to compress regular files. In particular, it will ignore symbolic links.

325      When compressing, gzip uses the deflate algorithm specified in RFC 1951: DEFLATE  
 326      Compressed Data Format Specification and stores the result in a file using the gzip  
 327      file format specified in RFC 1952: GZIP File Format Specification.

**Options**

328      -c, --stdout, --to-stdout

329            writes output on standard output, leaving the original files unchanged. If there  
 330            are several input files, the output consists of a sequence of independently  
 331            compressed members. To obtain better compression, concatenate all input files  
 332            before compressing them.

333      -d, --decompress, --uncompress

334            the name operands are compressed files, and **gzip** shall decompress them.

335      -f, --force

336            forces compression or decompression even if the file has multiple links or the  
 337            corresponding file already exists, or if the compressed data is read from or  
 338            written to a terminal. If the input data is not in a format recognized by **gzip**, and  
 339            if the option --stdout is also given, copy the input data without change to the  
 340            standard ouput: let **gzip** behave as **cat**. If -f is not given, and when not running  
 341            in the background, **gzip** prompts to verify whether an existing file should be  
 342            overwritten.

343      -l, --list

344            lists the compressed size, uncompressed size, ratio and uncompressed name for  
 345            each compressed file. For files that are not in **gzip** format, the uncompressed  
 346            size shall be given as -1. If the --verbose or -v option is also specified, the crc  
 347            and timestamp for the uncompressed file shall also be displayed.

348            For decompression, **gzip** shall support at least the following compression  
 349            methods:

- deflate (RFC 1951: DEFLATE Compressed Data Format Specification)
- compress (ISO POSIX (2003))

352            The crc shall be given as ffffffff for a file not in **gzip** format.

353        If the `--name` or `-N` option is also specified, the uncompressed name, date and  
 354        time are those stored within the compressed file, if present.

355        If the `--quiet` or `-q` option is also specified, the title and totals lines are not  
 356        displayed.

357     **-L, --license**  
 358        displays the **gzip** license and quit.

359     **-n, --no-name**  
 360        does not save the original file name and time stamp by default when  
 361        compressing. (The original name is always saved if the name had to be  
 362        truncated.) When decompressing, do not restore the original file name if present  
 363        (remove only the gzip suffix from the compressed file name) and do not restore  
 364        the original time stamp if present (copy it from the compressed file). This option  
 365        is the default when decompressing.

366     **-N, --name**  
 367        always saves the original file name and time stamp when compressing; this is  
 368        the default. When decompressing, restore the original file name and time stamp  
 369        if present. This option is useful on systems which have a limit on file name  
 370        length or when the time stamp has been lost after a file transfer.

371     **-q, --quiet**  
 372        suppresses all warnings.

373     **-r, --recursive**  
 374        travels the directory structure recursively. If any of the file names specified on  
 375        the command line are directories, **gzip** will descend into the directory and  
 376        compress all the files it finds there (or decompress them in the case of **gunzip**).

377     **-S .suf, --sufix .suf**  
 378        uses suffix `.suf` instead of `.gz`.

379     **-t, --test**  
 380        checks the compressed file integrity.

381     **-v, --verbose**  
 382        displays the name and percentage reduction for each file compressed or  
 383        decompressed.

384     **-#, --fast, --best**  
 385        regulates the speed of compression using the specified digit `#`, where `-1` or  
 386        `--fast` indicates the fastest compression method (less compression) and `-9` or  
 387        `--best` indicates the slowest compression method (best compression). The  
 388        default compression level is `-6` (that is, biased towards high compression at  
 389        expense of speed).

## LSB Deprecated Options

390        The behaviors specified in this section are expected to disappear from a future  
 391        version of the LSB; applications should only use the non-LSB-deprecated behaviors.

392	-V, --version
393	displays the version number and compilation options, then quits.
<b>hostname</b>	
<b>Name</b>	
394	hostname — show or set the system's host name
<b>Synopsis</b>	
395	<b>hostname</b> [name]
<b>Description</b>	
396	<b>hostname</b> is used to either display or, with appropriate privileges, set the current
397	host name of the system. The host name is used by many applications to identify the
398	machine.
399	When called without any arguments, the program displays the name of the system
400	as returned by the <code>gethostname()</code> function.
401	When called with a <i>name</i> argument, and the user has appropriate privilege, the
402	command sets the host name.
403	<b>Note:</b> It is not specified if the hostname displayed will be a fully qualified domain name.
404	Applications requiring a particular format of hostname should check the output and take
405	appropriate action.

## install

### Name

406      **install** – copy files and set attributes

### Synopsis

407      **install** [option...] SOURCE DEST **install** [option...] SOURCE... DEST **install** [-d  
408      | --directory] [option...] DIRECTORY...

### Description

409      In the first two formats, copy *SOURCE* to *DEST* or multiple *SOURCE(s)* to the existing  
410      *DEST* directory, optionally setting permission modes and file ownership. In the third  
411      format, each *DIRECTORY* and any missing parent directories shall be created.

### Standard Options

412      **--backup[=METHOD]**  
413          makes a backup of each existing destination file. *METHOD* may be one of the  
414          following:  
  
415          *none* or *off*  
416          never make backups.  
  
417          *numbered* or *t*  
418          make numbered backups. A numbered backup has the form "%s.~%d~" ,  
419          *target\_name*, *version\_number*. Each backup shall increment the version  
420          number by 1.  
  
421          *existing* or *nil*  
422          behave as numbered if numbered backups exist, or simple otherwise.  
  
423          *simple* or *never*  
424          append a suffix to the name. The default suffix is '~', but can be overridden  
425          by setting SIMPLE\_BACKUP\_SUFFIX in the environment, or via the *-s* or  
426          *--suffix* option.  
  
427      If no *METHOD* is specified, the environment variable VERSION\_CONTROL shall  
428      be examined for one of the above. Unambiguous abbreviations of *METHOD* shall  
429      be accepted. If no *METHOD* is specified, or if *METHOD* is empty, the backup method  
430      shall default to *existing*.  
  
431      If *METHOD* is invalid or ambiguous, **install** shall fail and issue a diagnostic  
432      message.  
  
433      **-b**  
434          is equivalent to *--backup=existing*.  
  
435      **-d, --directory**  
436          treats all arguments as directory names; creates all components of the specified  
437          directories.

438 -D  
439       creates all leading components of DEST except the last, then copies SOURCE to  
440       DEST; useful in the 1st format.

441 -g GROUP, --group=GROUP  
442       if the user has appropriate privilege, sets group ownership, instead of process'  
443       current group. *GROUP* is either a name in the user group database, or a positive  
444       integer, which shall be used as a group-id.

445 -m MODE, --mode=MODE  
446       sets permission mode (specified as in **chmod**), instead of the default rwxr-xr-x.

447 -o OWNER, --owner=OWNER  
448       if the user has appropriate privilege, sets ownership. *OWNER* is either a name in  
449       the user login database, or a positive integer, which shall be used as a user-id.

450 -p, --preserve-timestamps  
451       copies the access and modification times of *SOURCE* files to corresponding  
452       destination files.

453 -s, --strip  
454       strips symbol tables, only for 1st and 2nd formats.

455 -S SUFFIX, --suffix=SUFFIX  
456       equivalent to --backup=existing, except if a simple suffix is required, use  
457       *SUFFIX*.

458 --verbose  
459       prints the name of each directory as it is created.

460 -v, --verbose  
461       print the name of each file before copying it to `stdout`.

## install\_initd

### Name

462 `install_initd` — activate an init script

### Synopsis

463 `/usr/lib/lsb/install_initd initd_file`

### Description

464 `install_initd` shall activate a system initialization file that has been copied to an  
465 implementation defined location such that this file shall be run at the appropriate  
466 point during system initialization. The `install_initd` command is typically called in  
467 the postinstall script of a package, after the script has been copied to `/etc/init.d`.  
468 See also Installation and Removal of Init Scripts.

## ipcrm

### Name

469 ipcrm — remove IPC Resources

### Synopsis

470 ipcrm [-q msgid | -Q msgkey | -s semid | -S semkey | -m shmid | -M shmkey]... ipcrm  
 471 [shm | msg | msg] id...

### Description

472 If any of the `-q`, `-Q`, `-s`, `-S`, `-m`, or `-M` arguments are given, the **ipcrm** shall behave as  
 473 described in ISO POSIX (2003).

474 Otherwise, **ipcrm** shall remove the resource of the specified type identified by *id*.

### Future Directions

475 A future revision of this specification may deprecate the second synopsis form.

476 **Rationale:** In its first Linux implementation, **ipcrm** used the second syntax shown in the  
 477 SYNOPSIS. Functionality present in other implementations of **ipcrm** has since been  
 478 added, namely the ability to delete resources by key (not just identifier), and to respect  
 479 the same command line syntax. The previous syntax is still supported for backwards  
 480 compatibility only.

**ipcs****Name**

481      **ipcs** — provide information on ipc facilities

**Synopsis**

482      **ipcs** [-smq] [-tcp]

**Description**

483      **ipcs** provides information on the ipc facilities for which the calling process has read  
484      access.

485      **Note:** Although this command has many similarities with the optional **ipcs** utility  
486      described in ISO POSIX (2003), it has substantial differences and is therefore described  
487      separately. The options specified here have similar meaning to those in ISO POSIX  
488      (2003); other options specified there have unspecified behavior on an LSB conforming  
489      implementation. See Application Usage below. The output format is not specified.

**Resource display options**

490      -m  
491                shared memory segments.

492      -q  
493                message queues.

494      -s  
495                semaphore arrays.

**Output format options**

496      -t  
497                time.

498      -p  
499                pid.

500      -c  
501                creator.

**Application Usage**

502      In some implementations of ipcs the **-a** option will print all information available. In  
503      other implementations the **-a** option will print all resource types. Therefore,  
504      applications shall not use the **-a** option.

505      Some implementations of **ipcs** provide more output formats than are specified here.  
506      These options are not consistent between differing implementations of **ipcs**.  
507      Therefore, only the **-t**, **-c** and **-p** option formatting flags may be used. At least one  
508      of the **-t**, **-c** and **-p** options and at least one of **-m**, **-q** and **-s** options shall be  
509      specified. If no options are specified, the output is unspecified.

**killall****Name**

510      **killall** – kill processes by name

**Synopsis**

511      **killall** [-egiqvw] [-signal] name... **killall** -l **killall** -v

**Description**

512      **killall** sends a signal to all processes running any of the specified commands. If no  
513      signal name is specified, SIGTERM is sent.

514      Signals can be specified either by name (e.g. -HUP) or by number (e.g. -1). Signal 0  
515      (check if a process exists) can only be specified by number.

516      If the command name contains a slash (/), processes executing that particular file  
517      will be selected for killing, independent of their name.

518      **killall** returns a non-zero return code if no process has been killed for any of the  
519      listed commands. If at least one process has been killed for each command, **killall**  
520      returns zero.

521      A **killall** process never kills itself (but may kill other **killall** processes).

**Standard Options**

522      -e

523      requires an exact match for very long names. If a command name is longer than  
524      15 characters, the full name may be unavailable (i.e. it is swapped out). In this  
525      case, **killall** will kill everything that matches within the first 15 characters. With  
526      -e, such entries are skipped. **killall** prints a message for each skipped entry if -v  
527      is specified in addition to -e.

528      -g

529      kills the process group to which the process belongs. The kill signal is only sent  
530      once per group, even if multiple processes belonging to the same process group  
531      were found.

532      -i

533      asks interactively for confirmation before killing.

534      -l

535      lists all known signal names.

536      -q

537      does not complain if no processes were killed.

538      -v

539      reports if the signal was successfully sent.

**LSB Deprecated Options**

540        The behaviors specified in this section are expected to disappear from a future  
541        version of the LSB; applications should only use the non-LSB-deprecated behaviors.

542        -V  
543              displays version information.

## **lpr**

### **Name**

544        lpr – off line print

### **Synopsis**

545        **lpr** [-l] [-p] [-Pprinter] [-h] [-s] [-#copies] [-J name] [-T title] [name .....]

### **Description**

546        lpr uses a spooling daemon to print the named files when facilities become available.  
547        If no names appear, the standard input is assumed.

### **Standard Options**

548        -l  
549              identifies binary data that is not to be filtered but sent as raw input to printer.  
  
550        -p  
551              formats with "pr" before sending to printer.  
  
552        -Pprinter  
553              sends output to the printer named printer instead of the default printer.  
  
554        -h  
555              suppresses header page.  
  
556        -s  
557              uses symbolic links.  
  
558        -#copies  
559              specifies copies as the number of copies to print.  
  
560        -J name  
561              specifies name as the job name for the header page.  
  
562        -T title  
563              specifies title as the title used for "pr".

**ls****Name**

564        **ls** – list directory contents

**Description**

565        **ls** shall behave as specified in ISO POSIX (2003), but with extensions listed below.

**Extensions**

566        **-l**

567        If the file is a character special or block special file, the size of the file shall be  
568        replaced with two unsigned numbers in the format "%u, %u", representing the  
569        major and minor device numbers associated with the special file.

570        **Note:** The LSB does not specify the meaning of the major and minor devices numbers.

571        **-p**

572        in addition to ISO POSIX (2003) XSI optional behavior of printing a slash for a  
573        directory, **ls -p** may display other characters for other file types.

## lsb\_release

### Name

574        `lsb_release` – print distribution specific information

### Synopsis

575        `lsb_release [OPTION...]`

### Description

576        The `lsb_release` command prints certain LSB (Linux Standard Base) and  
577        Distribution information.

578        If no options are given, the `-v` option is assumed.

### Options

579        `-v, --version`

580        displays version of LSB against which distribution is compliant. The version is  
581        expressed as a colon separated list of LSB module descriptions. LSB module  
582        descriptions are dash separated tuples containing the module name, version,  
583        and architecture name. The output is a single line of text of the following  
584        format:

585        `LSB Version:\tListAsDescribedAbove`

586        **Note:** An implementation may support multiple releases of the same module.  
587        Version specific library interfaces, if any, will be selected by the program interpreter,  
588        which changes from release to release. Version specific commands and utilities, if  
589        any, will be described in the relevant specification.

590        `-i, --id`

591        displays string id of distributor. The output is a single line of text of the  
592        following format:

593        `Distributor ID:\tDistributorID`

594        `-d, --description`

595        displays single line text description of distribution. The output is of the  
596        following format:

597        `Description:\tDescription`

598        `-r, --release`

599        displays release number of distribution. The output is a single line of text of the  
600        following format:

601        `Release:\tRelease`

602        `-c, --codename`

603        displays codename according to distribution release. The output is a single line  
604        of text of the following format:

605        `Codename:\tCodename`

606        `-a, --all`

607           displays all of the above information.  
608        -s, --short  
609           displays all of the above information in short output format.  
610        -h, --help  
611           displays a human-readable help message.

## Examples

612       The following command will list the LSB Profiles which are currently supported on  
613       this platform.

```
614 example% lsb_release -v  
615 LSB Version:  
616 core-3.1-ia32:core-3.1-noarch:graphics-3.1-ia32:graphics-3.1-noarch
```

# m4

## Name

617       m4 — macro processor

## Description

618       m4 is as specified in ISO POSIX (2003), but with extensions as listed below.

## Extensions

619       -P  
620           forces all builtins to be prefixed with m4\_. For example, define becomes  
621           m4\_define.  
622       -I *directory*  
623           Add *directory* to the end of the search path for includes.

## **md5sum**

### **Name**

624      **md5sum** — generate or check MD5 message digests

### **Synopsis**

625      **md5sum** [-c [file] | file]

### **Description**

626      For each file, write to standard output a line containing the MD5 message digest of  
627      that file, followed by one or more blank characters, followed by the name of the file.  
628      The MD5 message digest shall be calculated according to RFC 1321: The MD5  
629      Message-Digest Algorithm and output as 32 hexadecimal digits.

630      If no file names are specified as operands, read from standard input and use "-" as  
631      the file name in the output.

### **Options**

632      -c [file]

633            checks the MD5 message digest of all files named in *file* against the message  
634            digest listed in the same file. The actual format of *file* is the same as the output  
635            of **md5sum**. That is, each line in the file describes a file. If *file* is not specified,  
636            read message digests from `stdin`.

### **Exit Status**

637      **md5sum** shall exit with status 0 if the sum was generated successfully, or, in check  
638      mode, if the check matched. Otherwise, **md5sum** shall exit with a non-zero status.

## mknod

### Name

639 mknod — make special files

### Synopsis

640 **mknod** [-m mode | --mode=mode] name type [major minor]**mknod** [--version]

### Description

641 The **mknod** command shall create a special file named *name* of the given *type*.

642 The *type* shall be one of the following:

643 b

644 creates a block (buffered) special file with the specified *major* and *minor* device  
645 numbers.

646 c, u

647 creates a character (unbuffered) special file with the specified *major* and *minor*  
648 device numbers.

649 p

650 creates a FIFO.

### Options

651 -m mode, --mode=mode

652 create the special file with file access permissions set as described in *mode*. The  
653 permissions may be any absolute value (i.e. one not containing '+' or '-')  
654 acceptable to the **chmod** command.

655 --version

656 output version information and exit.

657 **Note:** This option may be deprecated in a future release of this specification.

658 If *type* is p, *major* and *minor* shall not be specified. Otherwise, these parameters are  
659 mandatory.

### Future Directions

660 This command may be deprecated in a future version of this specification. The  
661 *major* and *minor* operands are insufficiently portable to be specified usefully here.  
662 Only a FIFO can be portably created by this command, and the **mkfifo** command is a  
663 simpler interface for that purpose.

## **mktemp**

### **Name**

664      **mktemp** — make temporary file name (unique)

### **Synopsis**

665      **mktemp** [-q] [-u] *template*

### **Description**

666      The **mktemp** command takes the given file name *template* and overwrites a portion  
 667      of it to create a file name. This file name shall be unique and suitable for use by the  
 668      application.

669      The *template* should have at least six trailing 'x' characters. These characters are  
 670      replaced with characters from the portable filename character set in order to  
 671      generate a unique name.

672      If **mktemp** can successfully generate a unique file name, and the *-u* option is not  
 673      present, the file shall be created with read and write permission only for the current  
 674      user. The **mktemp** command shall write the filename generated to the standard  
 675      output.

### **Options**

676      **-q**

677            fail silently if an error occurs. Diagnostic messages to `stderr` are suppressed,  
 678            but the command shall still exit with a non-zero exit status if an error occurs.

679      **-u**

680            operates in `unsafe' mode. A unique name is generated, but the temporary file  
 681            shall be unlinked before **mktemp** exits. Use of this option is not encouraged.

**more****Name**

682      **more** — display files on a page-by-page basis

**Description**

683      **more** is as specified in ISO POSIX (2003), but with differences as listed below.

**Differences**

684      The **more** command need not respect the `LINES` and `COLUMNS` environment variables.

685      The following additional options may be supported:

686      `-num`

687            specifies an integer which is the screen size (in lines).

688      `+num`

689            starts at line number *num*.

690      `+/pattern`

691            Start at the first line matching the pattern, equivalent to executing the search  
692            forward (/) command with the given pattern immediately after opening each  
693            file.

694      The following options from ISO POSIX (2003) may behave differently:

695      `-e`

696            has unspecified behavior.

697      `-i`

698            has unspecified behavior.

699      `-n`

700            has unspecified behavior.

701      `-p`

702            Either clear the whole screen before displaying any text (instead of the usual  
703            scrolling behavior), or provide the behavior specified by ISO POSIX (2003). In  
704            the latter case, the syntax is "`-p command`".

705      `-t`

706            has unspecified behavior.

707      The **more** command need not support the following interactive commands:

```
g
G
u
control u
control f
newline
j
k
r
R
m
'(return to mark)
/!
?
N
:e
:t
control g
ZZ
```

708

## Rationale

709 The *+num* and *+/string* options are deprecated in SUSv2, and have been removed  
710 in ISO POSIX (2003); however this specification continues to specify them because  
711 the publicly available util-linux package does not support the replacement (-p  
712 *command*). The *+command* option as found in SUSv2 is more general than is specified  
713 here, but the util-linux package appears to only support the more specific *+num*  
714 and *+/string* forms.

## mount

### Name

715 mount — mount a file system

### Synopsis

716 **mount** [-hV]**mount** [-a] [-fFnrvw] [-t *vfstype*]**mount** [-fnrvw] [-o *options* [ , . . . ]]  
 717 [device | dir]**mount** [-fnrvw] [-t *vfstype*] [-o *options*] device dir

### Description

718 As described in ISO POSIX (2003), all files in the system are organized in a directed  
 719 graph, known as the file hierarchy, rooted at /. These files can be spread out over  
 720 several underlying devices. The **mount** command shall attach the file system found  
 721 on some underlying device to the file hierarchy.

### Options

722 -v  
 723       invoke verbose mode. The **mount** command shall provide diagnostic messages  
 724       on stdout.  
 725 -a  
 726       mount all file systems (of the given types) mentioned in /etc/fstab.  
 727 -F  
 728       If the -a option is also present, fork a new incarnation of **mount** for each device  
 729       to be mounted. This will do the mounts on different devices or different NFS  
 730       servers in parallel.  
 731 -f  
 732       cause everything to be done except for the actual system call; if it's not obvious,  
 733       this 'fakes' mounting the file system.  
 734 -n  
 735       mount without writing in /etc/mtab. This is necessary for example when /etc  
 736       is on a read-only file system.  
 737 -s  
 738       ignore **mount** options not supported by a file system type. Not all file systems  
 739       support this option.  
 740 -r  
 741       mount the file system read-only. A synonym is -o ro.  
 742 -w  
 743       mount the file system read/write. (default) A synonym is -o rw.  
 744 -L label

745                    If the file /proc/partitions is supported, mount the partition that has the  
746                    specified label.

747     -U *uuid*  
748                    If the file /proc/partitions is supported, mount the partition that has the  
749                    specified *uuid*.

750     -t *vfstype*  
751                    indicate a file system type of *vfstype*.  
752                    More than one type may be specified in a comma separated list. The list of file  
753                    system types can be prefixed with no to specify the file system types on which  
754                    no action should be taken.

755     -o  
756                    options are specified with a -o flag followed by a comma-separated string of  
757                    options. Some of these options are only useful when they appear in the  
758                    /etc/fstab file. The following options apply to any file system that is being  
759                    mounted:

760        async  
761                    perform all I/O to the file system asynchronously.

762        atime  
763                    update inode access time for each access. (default)

764        auto  
765                    in /etc/fstab, indicate the device is mountable with -a.

766        defaults  
767                    use default options: rw, suid, dev, exec, auto, nouser, async.

768        dev  
769                    interpret character or block special devices on the file system.

770        exec  
771                    permit execution of binaries.

772        noatime  
773                    do not update file access times on this file system.

774        noauto  
775                    in /etc/fstab, indicates the device is only explicitly mountable.

776        nodev  
777                    do not interpret character or block special devices on the file system.

778        noexec  
779                    do not allow execution of any binaries on the mounted file system.

780        nosuid

781                    do not allow set-user-identifier or set-group-identifier bits to take effect.

782        **nouser**

783                    forbid an unprivileged user to mount the file system. (default)

784        **remount**

785                    remount an already-mounted file system. This is commonly used to change  
786                    the mount options for a file system, especially to make a read-only file  
787                    system writable.

788        **ro**

789                    mount the file system read-only.

790        **rw**

791                    mount the file system read-write.

792        **suid**

793                    allow set-user-identifier or set-group-identifier bits to take effect.

794        **sync**

795                    do all I/O to the file system synchronously.

796        **user**

797                    allow an unprivileged user to mount the file system. This option implies  
798                    the options `noexec`, `nosuid`, `nodev` unless overridden by subsequent  
799                    options.

## LSB Deprecated Options

The behaviors specified in this section are expected to disappear from a future version of the LSB; applications should only use the non-LSB-deprecated behaviors.

**-V**

output version and exit.

## msgfmt

### Name

msgfmt — create a message object from a message file

### Synopsis

804   **msgfmt** [options...] *filename*...

### Description

806   The **msgfmt** command generates a binary message catalog from a textual translation  
807   description. Message catalogs, or message object files, are stored in files with a **.mo**  
808   extension.

809   **Note:** The format of message object files is not guaranteed to be portable. Message  
810   catalogs should always be generated on the target architecture using the **msgfmt**  
811   command.

812   The source message files, otherwise known as portable object files, have a **.po**  
813   extension.

814   The *filename* operands shall be portable object files. The **.po** file contains messages  
815   to be displayed to users by system utilities or by application programs. The portable  
816   object files are text files, and the messages in them can be rewritten in any language  
817   supported by the system.

818   If any *filename* is **-**, a portable object file shall be read from the standard input.

819   The **msgfmt** command interprets data as characters according to the current setting  
820   of the **LC\_CTYPE** locale category.

### Options

821   **-c**

822   **--check**

823   Detect and diagnose input file anomalies which might represent translation  
824   errors. The **msgid** and **msgstr** strings are studied and compared. It is  
825   considered abnormal that one string starts or ends with a newline while the  
826   other does not.

827   If the message is flagged as **c-format** (see Comment Handling), check that the  
828   **msgid** string and the **msgstr** translation have the same number of % format  
829   specifiers, with matching types.

830   **-D** *directory*

831   **--directory=***directory*

832   Add *directory* to list for input files search. If *filename* is not an absolute  
833   pathname and *filename* cannot be opened, search for it in *directory*. This  
834   option may be repeated. Directories shall be searched in order, with the leftmost  
835   *directory* searched first.

836   **-f**

837   **--use-fuzzy**

838   Use entries marked as **fuzzy** in output. If this option is not specified, such  
839   entries are not included into the output. See Comment Handling below.

```

840 -o output-file
841 --output-file=output-file
842     Specify the output file name as output-file. If multiple domains or duplicate
843     msgids in the .po file are present, the behavior is unspecified. If output-file is -,
844     output is written to standard output.

845 --strict
846     Ensure that all output files have a .mo extension. Output files are named either
847     by the -o (or --output-file) option, or by domains found in the input files.

848 -v
849 --verbose
850     Print additional information to the standard error, including the number of
851     translated strings processed.

```

## Operands

852 The *filename* operands are treated as portable object files. The format of portable
853 object files is defined in EXTENDED DESCRIPTION.

## Standard Input

854 The standard input is not used unless a *filename* operand is specified as "-".

## Environment Variables

```

855 LANGUAGE
856     Specifies one or more locale names.

857 LANG
858     Specifies locale name.

859 LC_ALL
860     Specifies locale name for all categories. If defined, overrides LANG, LC_CTYPE
861     and LC_MESSAGES.

862 LC_CTYPE
863     Determine the locale for the interpretation of sequences of bytes of text data as
864     characters (for example, single-byte as opposed to multi-byte characters in
865     arguments and input files).

866 LC_MESSAGES
867     Specifies messaging locale, and if present overrides LANG for messages.

```

## Standard Output

868 The standard output is not used unless the option-argument of the -o option is
869 specified as -.

## Extended Description

870       The format of portable object files (.po files) is defined as follows. Each .po file  
 871       contains one or more lines, with each line containing either a comment or a  
 872       statement. Comments start the line with a hash mark (#) and end with the newline  
 873       character. Empty lines, or lines containing only white-space, shall be ignored.  
 874       Comments can in certain circumstances alter the behavior of **msgfmt**. See Comment  
 875       Handling below for details on comment processing. The format of a statement is:

876       directive value

877       Each directive starts at the beginning of the line and is separated from value by  
 878       white space (such as one or more space or tab characters). The value consists of one  
 879       or more quoted strings separated by white space. If two or more strings are specified  
 880       as value, they are normalized into single string using the string normalization  
 881       syntax specified in ISO C (1999). The following directives are supported:

882       domain domainname

883       msgid message\_identifier

884       msgid\_plural untranslated\_string\_plural

885       msgstr message\_string

886       msgstr[n] message\_string

887       The behavior of the domain directive is affected by the options used. See OPTIONS  
 888       for the behavior when the -o option is specified. If the -o option is not specified, the  
 889       behavior of the domain directive is as follows:

- 890       1. All msgids from the beginning of each .po file to the first domain directive are  
        891       put into a default message object file, messages (or messages.mo if the  
        892       --strict option is specified).
- 893       2. When **msgfmt** encounters a domain domainname directive in the .po file, all  
        894       following msgids until the next domain directive are put into the message  
        895       object file domainname (or domainname.mo if --strict option is specified).
- 896       3. Duplicate msgids are defined in the scope of each domain. That is, a msgid is  
        897       considered a duplicate only if the identical msgid exists in the same domain.
- 898       4. All duplicate msgids are ignored.

899       The msgid directive specifies the value of a message identifier associated with the  
 900       directive that follows it. The msgid\_plural directive specifies the plural form  
 901       message specified to the plural message handling functions `ngettext()`,  
 902       `dngettext()` or `dcngettext()`. The message\_identifier string identifies a target  
 903       string to be used at retrieval time. Each statement containing a msgid directive shall  
 904       be followed by a statement containing a msgstr directive or msgstr[n] directives.

905       The msgstr directive specifies the target string associated with the  
 906       message\_identifier string declared in the immediately preceding msgid directive.

907       The msgstr[n] (where n = 0, 1, 2, ...) directive specifies the target string to be used  
 908       with plural form handling functions `ngettext()`, `dngettext()` and `dncngettext()`.

909       Message strings can contain the following escape sequences:

910       **Table 15-1 Escape Sequences**

\n	newline
\t	tab
\v	vertical tab

\b	backspace
\r	carriage return
\f	formfeed
\\\	backslash
\ "	double quote
\ddd	octal bit pattern
\xHH	hexadecimal bit pattern

911

## Comment Handling

913  
914  
915

Comments are introduced by a #, and continue to the end of the line. The second character (i.e. the character following the #) has special meaning. Regular comments should follow a space character. Other comment types include:

916  
917  
918  
919  
920  
921

```
# normal-comments
#. automatic-comments
#: reference...
#, flag
```

Automatic and reference comments are typically generated by external utilities, and are not specified by the LSB. The **msgfmt** command shall ignore such comments.

922  
923  
924

**Note:** Portable object files may be produced by unspecified tools. Some of the comment types described here may arise from the use of such tools. It is beyond the scope of this specification to describe these tools.

925  
926

The #, comments require one or more flags separated by the comma (,) character. The following flags can be specified:

927

fuzzy

928  
929  
930  
931

This flag shows that the following `msgstr` string might not be a correct translation. Only the translator (i.e. the individual undertaking the translation) can judge if the translation requires further modification, or is acceptable as is. Once satisfied with the translation, the translator then removes this fuzzy flag.

932  
933  
934

If this flag is specified, the **msgfmt** utility will not generate the entry for the immediately following `msgid` in the output message catalog, unless the `--use-fuzzy` is specified.

935  
936

c-format  
no-c-format

937  
938  
939

The c-format flag indicates that the `msgid` string is used as format string by `printf()`-like functions. If the c-format flag is given for a string the **msgfmt** utility may perform additional tests to check the validity of the translation.

940

## Plurals

941       The msgid entry with empty string ("") is called the header entry and is treated  
 942       specially. If the message string for the header entry contains nplurals=value, the  
 943       value indicates the number of plural forms. For example, if nplurals=4, there are 4  
 944       plural forms. If nplurals is defined, there should be a plural=expression on the  
 945       same line, separated by a semicolon (;) character. The expression is a C language  
 946       expression to determine which version of msgstr[n] to be used based on the value  
 947       of n, the last argument of gettext(), dgettext() or dcgettext(). For example:

948       nplurals=2; plural=n == 1 ? 0 : 1

949       indicates that there are 2 plural forms in the language; msgstr[0] is used if n == 1,  
 950       otherwise msgstr[1] is used. Another example:

951       nplurals=3; plural=n==1 ? 0 : n==2 ? 1 : 2

952       indicates that there are 3 plural forms in the language; msgstr[0] is used if n == 1,  
 953       msgstr[1] is used if n == 2, otherwise msgstr[2] is used.

954       If the header entry contains charset=codeset string, the codeset is used to indicate  
 955       the codeset to be used to encode the message strings. If the output string's codeset is  
 956       different from the message string's codeset, codeset conversion from the message  
 957       strings's codeset to the output string's codeset will be performed upon the call of  
 958       gettext(), dgettext(), dcgettext(), gettext(), dgettext(), and  
 959       dcgettext(). The output string's codeset is determined by the current locale's  
 960       codeset (the return value of nl\_langinfo(CODESET)) by default, and can be changed  
 961       by the call of bind\_textdomain\_codeset().

## Exit Status

962       The following exit values are returned:

963       0

964           Successful completion.

965       >0

966           An error occurred.

## Application Usage

967       Neither msgfmt nor any gettext() function imposes a limit on the total length of a  
 968       message. Installing message catalogs under the C locale is pointless, since they are  
 969       ignored for the sake of efficiency.

## Examples

970       Example 1: Examples of creating message objects from message files.

971       In this example module1.po, module2.po and module3.po are portable message  
 972       object files.

```
973 example% cat module1.po
974
975 # default domain "messages"
976
977 msgid "message one"
978
979 msgstr "mensaje número uno"
980
```

```

981      #
982      domain "help_domain"
983
984      msgid "help two"
985
986      msgstr "ayuda número dos"
987
988      #
989      domain "error_domain"
990
991      msgid "error three"
992
993      msgstr "error número tres"
994
995
996

997 example% cat module2.po
998
999      # default domain "messages"
1000
1001      msgid "message four"
1002
1003      msgstr "mensaje número cuatro"
1004
1005      #
1006
1007      domain "error_domain"
1008
1009      msgid "error five"
1010
1011      msgstr "error número cinco"
1012
1013      #
1014
1015      domain "window_domain"
1016
1017      msgid "window six"
1018
1019      msgstr "ventana número seises"

1020 example% cat module3.po
1021
1022      # default domain "messages"
1023
1024      msgid "message seven"
1025
1026      msgstr "mensaje número siete"
1027

```

1028 The following command will produce the output files `messages`, `help_domain`, and  
1029 `error_domain`.

1030 `example% msgfmt module1.po`

1031 The following command will produce the output files `messages.mo`,  
1032 `help_domain.mo`, `error_domain.mo`, and `window_domain.mo`.

1033 `example% msgfmt module1.po module2.po`

1034 The following example will produce the output file `hello.mo`.

1035 `example% msgfmt -o hello.mo module3.po`

## **newgrp**

### **Name**

1036      newgrp — change group ID

### **Synopsis**

1037      **newgrp** [group]

### **Description**

1038      The **newgrp** command is as specified in ISO POSIX (2003), but with differences as  
1039      listed below.

### **Differences**

1040      The **-l** option specified in ISO POSIX (2003) need not be supported.  
1041

**od****Name**

1042        **od** — dump files in octal and other formats

**Synopsis**

1043        **od** [-abcdfilox] [-w *width* | --width=*width*] [-v] [-A *address\_base*] [-j *skip*] [-n *count*]  
 1044        [-t *type\_string*] [*file...*] **od** --traditional [*options*] [*file*] [[+]offset [.] [b]]  
 1045        [[+]*label* [.] [b]]

**Description**

1046        The **od** command shall provide all of the mandatory functionality specified in ISO  
 1047        POSIX (2003), but with extensions and differences to the XSI optional behavior as  
 1048        listed below.

**Extensions and Differences**

1049        -s

1050        unspecified behavior.

1051        **Note:** Applications wishing to achieve the ISO POSIX (2003) behavior for *-s* should  
 1052        instead use *-t d2*.

1053        -w*width*, --width[=*width*]

1054        each output line is limited to *width* bytes from the input.

1055        --traditional

1056        accepts arguments in traditional form, see Traditional Usage below.

1057        **Note:** The XSI optional behavior for offset handling described in ISO POSIX (2003) is not  
 1058        supported unless the *--traditional* option is also specified.

**Pre-POSIX and XSI Specifications**

1060        The LSB supports mixing options between the mandatory and XSI optional synopsis  
 1061        forms in ISO POSIX (2003). The LSB shall support the following options:

1062        -a

1063        is equivalent to *-t a*, selects named characters.

1064        -b

1065        is equivalent to *-t o1*, selects octal bytes.

1066        -c

1067        is equivalent to *-t c*, selects characters.

1068        -d

1069        is equivalent to *-t u2*, selects unsigned decimal two byte units.

1070        -f

1071        is equivalent to *-t fF*, selects floats.

1072           **-i**  
 1073           is equivalent to **-t d2**, selects decimal two byte units.  
**Note:** This usage may change in future releases; portable applications should use **-t d2**.

1075           **-l**  
 1076           is equivalent to **-t d4**, selects decimal longs.

1077           **-o**  
 1078           is equivalent to **-t o2**, selects octal two byte units.

1079           **-x**  
 1080           is equivalent to **-t x2**, selects hexadecimal two byte units.

1081 Note that the XSI option **-s** need not be supported.

1082 **Traditional Usage**

1083 If the **--traditional** option is specified, there may be between zero and three  
 1084 operands specified.

1085 If no operands are specified, then **od** shall read the standard input.

1086 If there is exactly one operand, and it is an offset of the form **[+]offset[.][b]**, then  
 1087 it shall be interpreted as specified in ISO POSIX (2003). The file to be dumped shall  
 1088 be the standard input.

1089 If there are exactly two operands, and they are both of the form **[+]offset[.][b]**,  
 1090 then the first shall be treated as an offset (as above), and the second shall be a label,  
 1091 in the same format as the offset. If a label is specified, then the first output line  
 1092 produced for each input block shall be preceded by the input offset, cumulative  
 1093 across input files, of the next byte to be written, followed by the label, in parentheses.  
 1094 The label shall increment in the same manner as the offset.

1095 If there are three operands, then the first shall be the file to dump, the second the  
 1096 offset, and the third the label.

1097 **Note:** Recent versions of **coreutils** contain an **od** utility that conforms to ISO POSIX  
 1098 (2003). However, in April 2005, this version was not in widespread use. A future version  
 1099 of this specification may remove the differences.

## **passwd**

### **Name**

1100      **passwd** — change user password

### **Synopsis**

1101    **passwd** [-x max] [-n min] [-w warn] [-i inact] name **passwd** {-l | -u} name

### **Description**

1102    **passwd** changes authentication information for user and group accounts, including  
 1103    passwords and password expiry details, and may be used to enable and disable  
 1104    accounts. Only a user with appropriate privilege may change the password for other  
 1105    users or modify the expiry information.

### **Options**

1106    **-x max**

1107         sets the maximum number of days a password remains valid.

1108    **-n min**

1109         sets the minimum number of days before a password may be changed.

1110    **-w warn**

1111         sets the number of days warning the user will receive before their password will  
 1112         expire.

1113    **-i inactive**

1114         disables an account after the password has been expired for the given number  
 1115         of days.

1116    **-l**

1117         disables an account by changing the password to a value which matches no  
 1118         possible encrypted value.

1119    **-u**

1120         re-enables an account by changing the password back to its previous value.

## patch

### Name

1121 patch — apply a diff file to an original

### Description

1122 patch is as specified in ISO POSIX (2003), but with extensions as listed below.

### Extensions

1123 --binary

1124 reads and write all files in binary mode, except for standard output and  
1125 /dev/tty. This option has no effect on POSIX-compliant systems.

1126 -u, --unified

1127 interprets the patch file as a unified context diff.

## pidof

### Name

1128 pidof — find the process ID of a running program

### Synopsis

1129 **pidof** [-s] [-x] [-o omitpid...] program...

### Description

1130 Return the process ID of a process which is running the program named on the  
1131 command line.

1132 The **pidof** command is a system administration utility, see Path For System  
1133 Administration Utilities.

### Options

1134 -s

1135 instructs the program to only return one pid.

1136 -x

1137 causes the program to also return process id's of shells running the named  
1138 scripts.

1139 -o

1140 omits processes with specified process id.

**remove\_initd****Name**

1141   **remove\_initd** — clean up init script system modifications introduced by  
 1142   **install\_initd**

**Synopsis**

1143   `/usr/lib/lsb/remove_initd initd_file`

**Description**

1144   **remove\_initd** processes the removal of the modifications made to a distribution's  
 1145   init script system by the **install\_initd** program. This cleanup is performed in the  
 1146   preuninstall script of a package; however, the package manager is still responsible  
 1147   for removing the script from the repository. See also Installation and Removal of Init  
 1148   Scripts.

**renice****Name**

1149   **renice** — alter priority of running processes

**Description**

1150   **renice** is as specified in ISO POSIX (2003), but with differences as listed below.

**Differences**

1151   -n increment  
 1152       has unspecified behavior.

**sed****Name**

1153   **sed** — stream editor

**Description**

1154   **sed** is as specified in ISO POSIX (2003), but with differences as listed below.

**LSB Differences**

1155   Certain aspects of internationalized regular expressions are optional; see  
 1156   Internationalization and Regular Expressions.

## sendmail

### Name

1157 sendmail – an electronic mail transport agent

### Synopsis

1158 `/usr/sbin/sendmail [options] [address...]`

### Description

1159 To deliver electronic mail (email), applications shall support the interface provided  
 1160 by **sendmail** (described here). This interface shall be the default delivery method for  
 1161 applications.

1162 This program sends an email message to one or more recipients, routing the message  
 1163 as necessary. This program is not intended as a user interface routine.

1164 With no options, **sendmail** reads its standard input up to an end-of-file or a line  
 1165 consisting only of a single dot and sends a copy of the message found there to all of  
 1166 the addresses listed. It determines the network(s) to use based on the syntax and  
 1167 contents of the addresses.

1168 If an address is preceded by a backslash, '\', it is unspecified if the address is  
 1169 subject to local alias expansion.

1170 The format of messages shall be as defined in RFC 2822:Internet Message Format.

1171 **Note:** The name **sendmail** was chosen for historical reasons, but the **sendmail** command  
 1172 specified here is intended to reflect functionality provided by **smail**, **exim** and other  
 1173 implementations, not just the **sendmail** implementation.

### Options

1174 -bm

1175 read mail from standard input and deliver it to the recipient addresses. This is  
 1176 the default mode of operation.

1177 -bp

1178 If the user has sufficient privilege, list information about messages currently in  
 1179 the mail queue.

1180 -bs

1181 use the SMTP protocol as described in RFC 2821:Simple Mail Transfer Protocol;  
 1182 read SMTP commands on standard input and write SMTP responses on  
 1183 standard output.

1184 In this mode, **sendmail** shall accept \r\n (CR-LF), as required by RFC  
 1185 2821:Simple Mail Transfer Protocol, and \n (LF) line terminators.

1186 -F fullname

1187 explicitly set the full name of the sender for incoming mail unless the message  
 1188 already contains a `From:` message header.

1189 If the user running **sendmail** is not sufficiently trusted, then the actual sender  
 1190 may be indicated in the message, depending on the configuration of the agent.

1191 -f name  
 1192     explicitly set the envelope sender address for incoming mail. If there is no `From:`  
 1193     header, the address specified in the `From:` header will also be set.  
 1194     If the user running **sendmail** is not sufficiently trusted, then the actual sender  
 1195     shall be indicated in the message.

1196 -i  
 1197     ignore dots alone on lines by themselves in incoming messages. If this options is  
 1198     not specified, a line consisting of a single dot shall terminate the input. If `-bs` is  
 1199     also used, the behavior is unspecified.

1200 -odb  
 1201     deliver any mail in background, if supported; otherwise ignored.

1202 -odf  
 1203     deliver any mail in foreground, if supported; otherwise ignored.

1204 -oem or -em  
 1205     mail errors back to the sender. (default)

1206 -oep or -ep  
 1207     write errors to the standard error output.

1208 -oeq or -eq  
 1209     do not send notification of errors to the sender. This only works for mail  
 1210     delivered locally.

1211 -oi  
 1212     is equivalent to `-i`.

1213 -om  
 1214     indicate that the sender of a message should receive a copy of the message if the  
 1215     sender appears in an alias expansion. Ignored if aliases are not supported.

1216 -t  
 1217     read the message to obtain recipients from the `To:`, `Cc:`, and `Bcc:` headers in the  
 1218     message instead of from the command arguments. If a `Bcc:` header is present, it  
 1219     is removed from the message unless there is no `To:` or `Cc:` header, in which case  
 1220     a `Bcc:` header with no data is created, in accordance with RFC 2822:Internet  
 1221     Message Format.  
 1222     If there are any operands, the recipients list is unspecified.  
 1223     This option may be ignored when not in `-bm` mode (the default).

1224     **Note:** It is recommended that applications use as few options as necessary, none if  
 1225     possible.

## Exit status

1226 0

1227                   successful completion on all addresses. This does not indicate successful  
1228                   delivery.

1229                   >0  
1230                   there was an error.

## sh

### Name

1231                   sh — shell, the standard command language interpreter

### Description

1232                   The **sh** utility shall behave as specified in ISO POSIX (2003), but with extensions  
1233                   listed below.

### Shell Invocation

1234                   The shell shall support an additional option, **-l** (the letter *ell*). If the **-l** option is  
1235                   specified, or if the first character of argument zero (the command name) is a '**-**', this  
1236                   invocation of the shell is a *login shell*.

1237                   An interactive shell, as specified in ISO POSIX (2003), that is also a login shell, or any  
1238                   shell if invoked with the **-l** option, shall, prior to reading from the input file, first  
1239                   read and execute commands from the file **/etc/profile**, if that file exists, and then  
1240                   from a file called **~/.profile**, if such a file exists.

1241                   **Note:** This specification requires that the **sh** utility shall also read and execute  
1242                   commands in its current execution environment from all the shell scripts in the directory  
1243                   **/etc/profile.d**. Such scripts are read and executed as a part of reading and executing  
1244                   **/etc/profile**.

## shutdown

### Name

1245 shutdown — shut the system down

### Synopsis

1246 /sbin/shutdown [-t sec] [-h | -r] [-akfF] time [warning-message]/sbin/shutdown  
1247 -c [warning-message]

### Description

1248 The **shutdown** command shall shut the system down in a secure way (first synopsis),  
1249 or cancel a pending shutdown (second synopsis). When the shutdown is initiated, all  
1250 logged-in users shall be notified immediately that the system is going down, and  
1251 users shall be prevented from logging in to the system. The *time* specifies when the  
1252 actual shutdown shall commence. See below for details. At the specified time all  
1253 processes are first notified that the system is going down by the signal SIGTERM.  
1254 After an interval (see *-t*) all processes shall be sent the signal SIGKILL. If neither the  
1255 *-h* or the *-r* argument is specified, then the default behavior shall be to take the  
1256 system to a runlevel where administrative tasks can be run. See also Run Levels.

1257 **Note:** This is sometimes referred to as "single user mode".

1258 The *-h* and *-r* options are mutually exclusive. If either the *-h* or *-r* options are  
1259 specified, the system shall be halted or rebooted respectively.

### Standard Options

1260 -a  
1261       use access control. See below.

1262 -t sec  
1263       tell the system to wait *sec* seconds between sending processes the warning and  
1264       the kill signal, before changing to another runlevel. The default period is  
1265       unspecified.

1266 -k  
1267       do not really shutdown; only send the warning messages to everybody.

1268 -r  
1269       reboot after shutdown.

1270 -h  
1271       halt after shutdown. Actions after halting are unspecified (e.g. power off).

1272 -f  
1273       advise the system to skip file system consistency checks on reboot.

1274 -F  
1275       advise the system to force file system consistency checks on reboot.

1276 -c

1277                   cancel an already running **shutdown**.

1278                   time

1279                   specify when to shut down.

1280                   The time argument shall have the following format: [now | [+]**mins** | hh:mm]

1281                   If the format is hh:mm, hh shall specify the hour (1 or 2 digits) and mm is the

1282                   minute of the hour (exactly two digits), and the shutdown shall commence at

1283                   the next occurrence of the specified time. If the format is **mins** (or +**mins**), where

1284                   **mins** is a decimal number, shutdown shall commence in the specified number

1285                   of minutes. The word now is an alias for +0.

1286                   warning-message

1287                   specify a message to send to all users.

1288                   **Access Control**

1289                   If the **shutdown** utility is invoked with the -a option, it shall check that an

1290                   authorized user is currently logged in on the system console. Authorized users are

1291                   listed, one per line, in the file /etc/shutdown.allow. Lines in this file that begin

1292                   with a '#' or are blank shall be ignored.

1293                   **Note:** The intent of this scheme is to allow a keyboard sequence entered on the system

1294                   console (e.g. CTRL-ALT-DEL, or STOP-A) to automatically invoke **shutdown -a**, and can be

1295                   used to prevent unauthorized users from shutting the system down in this fashion.

**su****Name**

1296    **su** — change user ID

**Synopsis**

1297    **su** [options] [-] [username [ARGS]]

**Description**

1298    The **su** command shall start a shell running with the real and effective user and  
 1299    group IDs of the user *username*. If *username* is not specified, **su** shall default to an  
 1300    unspecified user with all appropriate privileges. If the **-s** or **--shell** is not specified,  
 1301    the shell to be invoked shall be that specified for *username* in the user database (see  
 1302    `getpwnam()`), or `/bin/sh` if there is no shell specified in the user database.

1303    If the **-** option is specified, or if the first operand is **-**, the environment for the shell  
 1304    shall be initialized as if the new shell was a login shell (see Shell Invocation).

1305    If the invoking user does not have appropriate privileges, the **su** command shall  
 1306    prompt for a password and validate this before continuing. Invalid passwords shall  
 1307    produce an error message. The **su** command shall log in an unspecified manner all  
 1308    invocations, whether successful or unsuccessful.

1309    Any operands specified after the *username* shall be passed to the invoked shell.

1310    If the option **-** is not specified, and if the first operand is not **-**, the environment for  
 1311    the new shell shall be initialized from the current environment. If none of the **-m**, **-p**,  
 1312    or **--preserve-environment** options are specified, the environment may be  
 1313    modified in unspecified ways before invoking the shell. If any of the **-m**, **-p**, or  
 1314    **--preserve-environment** options are specified, the environment shall not be  
 1315    altered.

1316    **Note:** Although the **su** command shall not alter the environment, the invoked shell may  
 1317    still alter it before it is ready to interpret any commands.

**Standard Options**

1318    **-**

1319    the invoked shell shall be a login shell.

1320    **-c** *command*, **--command**=*command*

1321    Invoke the shell with the option **-c** *command*.

1322    **-m**, **-p**, **--preserve-environment**

1323    The current environment shall be passed to the invoked shell. If the  
 1324    environment variable `SHELL` is set, it shall specify the shell to invoke, if it  
 1325    matches an entry in `/etc/shells`. If there is no matching entry in `/etc/shells`,  
 1326    this option shall be ignored if the **-** option is also specified, or if the first  
 1327    operand is **-**.

1328    **-s** *shell*, **--shell**=*shell*

1329    Invoke *shell* as the command interpreter. The shell specified shall be present in  
 1330    `/etc/shells`.

## **sync**

### **Name**

1331      **sync** — flush file system buffers

### **Synopsis**

1332      **sync**

### **Description**

1333      Force changed blocks to disk, update the super block.

## **tar**

### **Name**

1334      **tar** — file archiver

### **Description**

1335      **tar** is as specified in SUSv2, but with differences as listed below.

### **Differences**

1336      Some elements of the Pattern Matching Notation are optional; see  
1337      Internationalization and Pattern Matching Notation.

1338      **-h**

1339      doesn't dump symlinks; dumps the files they point to.

1340      **-z**

1341      filters the archive through **gzip**.

## **umount**

### **Name**

1342      **umount** — unmount file systems

### **Synopsis**

1343    **umount** [-hV]**umount** -a [-nrv] [-t *vfstype*]**umount** [-nrv] *device* | *dir*

### **Description**

1344      **umount** detaches the file system(s) mentioned from the file hierarchy. A file system  
1345      is specified by giving the directory where it has been mounted.

### **Standard Options**

1346      -v

1347                invokes verbose mode.

1348      -n

1349                unmounts without writing in /etc/mtab.

1350      -r

1351                tries to remount read-only if unmounting fails.

1352      -a

1353                unmounts all of the file systems described in /etc/mtab except for the proc file  
1354                system.

1355      -t *vfstype*

1356                indicates that the actions should only be taken on file systems of the specified  
1357                type. More than one type may be specified in a comma separated list. The list of  
1358                file system types can be prefixed with no to specify the file system types on  
1359                which no action should be taken.

1360      -f

1361                forces unmount (in case of an unreachable NFS system).

### **LSB Deprecated Options**

1362      The behaviors specified in this section are expected to disappear from a future  
1363      version of the LSB; applications should only use the non-LSB-deprecated behaviors.

1364      -V

1365                print version and exits.

## **useradd**

### **Name**

1366   **useradd** — create a new user or update default new user information

### **Synopsis**

1367   **useradd** [-c comment] [-d home\_dir] [-g initial\_group] [-G group...] [-m [-k  
1368    skeleton\_dir]] [-p passwd] [-r] [-s shell] [-u uid [-o]] login **useradd** -D [-g  
1369    default\_group] [-b default\_home] [-s default\_shell]

### **Description**

1370   When invoked without the **-D** option, and with appropriate privilege, **useradd**  
1371   creates a new user account using the values specified on the command line and the  
1372   default values from the system. The new user account will be entered into the  
1373   system files as needed, the home directory will be created, and initial files copied,  
1374   depending on the command line options.

1375   When invoked with the **-D** option, **useradd** will either display the current default  
1376   values, or, with appropriate privilege, update the default values from the command  
1377   line. If no options are specified, **useradd** displays the current default values.

1378   The **useradd** command is a system administration utility, see Path For System  
1379   Administration Utilities.

### **Standard Options**

1380   **-c comment**

1381       specifies the new user's password file comment field value.

1382   **-d home\_dir**

1383       creates the new user using home\_dir as the value for the user's login directory.  
1384       The default is to append the login name to default\_home and use that as the  
1385       login directory name.

1386   **-g initial\_group**

1387       specifies the group name or number of the user's initial login group. The group  
1388       name shall exist. A group number shall refer to an already existing group. If **-g**  
1389       is not specified, the implementation will follow the normal user default for that  
1390       system. This may create a new group or choose a default group that normal  
1391       users are placed in. Applications which require control of the groups into which  
1392       a user is placed should specify **-g**.

1393   **-G group[...]**

1394       specifies a list of supplementary groups which the user is also a member of.  
1395       Each group is separated from the next by a comma, with no intervening  
1396       whitespace. The groups are subject to the same restrictions as the group given  
1397       with the **-g** option. The default is for the user to belong only to the initial group.

1398   **-m [-k skeleton\_dir]**

1399 specifies the user's home directory will be created if it does not exist. The files  
 1400 contained in `skeleton_dir` will be copied to the home directory if the `-k` option  
 1401 is used, otherwise the files contained in `/etc/skel` will be used instead. Any  
 1402 directories contained in `skeleton_dir` or `/etc/skel` will be created in the  
 1403 user's home directory as well. The `-k` option is only valid in conjunction with  
 1404 the `-m` option. The default is to not create the directory and to not copy any files.

1405 **-p passwd**

1406     is the encrypted password, as returned by `crypt()`. The default is to disable the  
 1407 account.

1408 **-r**

1409     creates a system account, that is, a user with a User ID in the range reserved for  
 1410 system account users. If there is not a User ID free in the reserved range the  
 1411 command will fail.

1412 **-s shell**

1413     specifies the name of the user's login shell. The default is to leave this field blank,  
 1414 which causes the system to select the default login shell.

1415 **-u uid [-o]**

1416     specifies the numerical value of the user's ID. This value shall be unique, unless  
 1417 the `-o` option is used. The value shall be non-negative. The default is the  
 1418 smallest ID value greater than 499 which is not yet used.

## Change Default Options

1419 **-b default\_home**

1420     specifies the initial path prefix for a new user's home directory. The user's name  
 1421 will be affixed to the end of `default_home` to create the new directory name if  
 1422 the `-d` option is not used when creating a new account.

1423 **-g default\_group**

1424     specifies the group name or ID for a new user's initial group. The named group  
 1425 shall exist, and a numerical group ID shall have an existing entry.

1426 **-s default\_shell**

1427     specifies the name of the new user's login shell. The named program will be  
 1428 used for all future new user accounts.

1429 **-c comment**

1430     specifies the new user's password file comment field value.

## Application Usage

1431 The `-D` option will typically be used by system administration packages. Most  
 1432 applications should not change defaults which will affect other applications and  
 1433 users.

## **userdel**

### **Name**

1434      **userdel** – delete a user account and related files

### **Synopsis**

1435      **userdel** [-r] *login*

### **Description**

1436      Delete the user account named *login*. If there is also a group named *login*, this  
1437      command may delete the group as well, or may leave it alone.

1438      The **userdel** command is a system administration utility, see Path For System  
1439      Administration Utilities.

### **Options**

1440      -r

1441      removes files in the user's home directory along with the home directory itself.  
1442      Files located in other file system will have to be searched for and deleted  
1443      manually.

## **usermod**

### **Name**

1444      **usermod** — modify a user account

### **Synopsis**

1445    **usermod** [-c comment] [-d home\_dir [-m]] [-g initial\_group] [-G group [,...]] [-l  
1446    login\_name] [-p passwd] [-s shell] [-u uid [-o]] login

### **Description**

1447    The **usermod** command shall modify an entry in the user account database.

1448    The **usermod** command is a system administration utility, see Path For System  
1449    Administration Utilities.

### **Options**

1450    -c comment

1451        specifies the new value of the user's password file comment field.

1452    -d home\_dir

1453        specifies the user's new login directory. If the -m option is given the contents of  
1454        the current home directory will be moved to the new home directory, which is  
1455        created if it does not already exist.

1456    -g initial\_group

1457        specifies the group name or number of the user's new initial login group. The  
1458        group name shall exist. A group number shall refer to an already existing  
1459        group.

1460    -G group[...]

1461        specifies a list of supplementary groups which the user is also a member of.  
1462        Each group is separated from the next by a comma, with no intervening  
1463        whitespace. The groups are subject to the same restrictions as the group given  
1464        with the -g option. If the user is currently a member of a group which is not  
1465        listed, the user will be removed from the group.

1466    -l login\_name

1467        changes the name of the user from login to login\_name. Nothing else is changed.  
1468        In particular, the user's home directory name should probably be changed to  
1469        reflect the new login name.

1470    -p passwd

1471        is the encrypted password, as returned by crypt(3).

1472    -s shell

1473        specifies the name of the user's new login shell. Setting this field to blank causes  
1474        the system to select the default login shell.

1475    -u uid [-o]

1476 specifies the numerical value of the user's ID. This value shall be unique, unless  
1477 the -o option is used. The value shall be non-negative. Any files which the user  
1478 owns and which are located in the directory tree rooted at the user's home  
1479 directory will have the file user ID changed automatically. Files outside of the  
1480 user's home directory shall be altered manually.

## xargs

### Name

1481 xargs — build and execute command lines from standard input

### Description

1482 xargs is as specified in ISO POSIX (2003), but with differences as listed below.

### Differences

1483 -E

1484 has unspecified behavior.

1485 -I

1486 has unspecified behavior.

1487 -L

1488 has unspecified behavior.

1489 **Note:** These options have been implemented in **findutils-4.2.9**, but this version of the  
1490 utilities is not in widespread use as of April 2005. However, future versions of this  
1491 specification will require support for these arguments.

## VI Execution Environment

## 16 File System Hierarchy

1 An LSB conforming implementation shall provide the mandatory portions of the file  
2 system hierarchy specified in the Filesystem Hierarchy Standard (FHS), together  
3 with any additional requirements made in this specification.

4 An LSB conforming application shall conform to the Filesystem Hierarchy Standard.

5 The FHS allows many components or subsystems to be optional. An application  
6 shall check for the existence of an optional component before using it, and should  
7 behave in a reasonable manner if the optional component is not present.

8 The FHS requirement to locate the operating system kernel in either / or /boot does  
9 not apply if the operating system kernel does not exist as a file in the file system.

10 The FHS specifies certain behaviors for a variety of commands if they are present  
11 (for example, **ping** or **python**). However, LSB conforming applications shall not rely  
12 on any commands beyond those specified by the LSB. The mere existence of a  
13 command may not be used as an indication that the command behaves in any  
14 particular way.

15 The following directories or links need not be present: /etc/X11 /usr/bin/X11  
16 /usr/lib/X11 /proc

### 16.1 /dev: Device Files

17 The following shall exist under /dev. Other devices may also exist in /dev. Device  
18 names may exist as symbolic links to other device nodes located in /dev or  
19 subdirectories of /dev. There is no requirement concerning major/minor number  
20 values.

21 /dev/null

22 An infinite data source and data sink. Data written to this device shall be  
23 discarded. Reads from this device shall always return end-of-file (EOF).

24 /dev/zero

25 This device is a source of zeroed out data. All data written to this device shall be  
26 discarded. A read from this device shall always return the requested number of  
27 bytes, each initialized to the value '\0'.

28 /dev/tty

29 In each process, a synonym for the controlling terminal associated with the  
30 process group of that process, if any. All reads and writes to this device shall  
31 behave as if the actual controlling terminal device had been opened.

### 16.2 /etc: Host-specific system configuration

32 In addition to the requirements for /etc in the Filesystem Hierarchy Standard, an  
33 LSB conforming system shall also provide the following directories or symbolic links  
34 to directories:

35 /etc/cron.d

36 A directory containing extended **crontab** files; see Cron Jobs.

37            /etc/cron.daily  
38            A directory containing shell scripts to be executed once a day; see Cron Jobs.  
39            /etc/cron.hourly  
40            A directory containing shell scripts to be executed once per hour; see Cron Jobs.  
41            /etc/cron.monthly  
42            A directory containing shell scripts to be executed once per month; see Cron  
43            Jobs.  
44            /etc/cron.weekly  
45            A directory containing shell scripts to be executed once a week; see Cron Jobs.  
46            /etc/init.d  
47            A directory containing system initialization scripts; see Installation and  
48            Removal of Init Scripts.  
49            /etc/profile.d  
50            A directory containing shell scripts. Script names should follow the same  
51            conventions as specified for cron jobs (see Cron Jobs, but should have the  
52            suffix .sh. The behavior is unspecified if a script is installed in this directory  
53            that does not have the suffix .sh.  
54            The **sh** utility shall read and execute commands in its current execution  
55            environment from all the shell scripts in this directory that have the suffix .sh  
56            when invoked as an interactive login shell, or if the -l (the letter *ell*) is specified  
57            (see Shell Invocation).  
58            **Future Directions:** These directories are required at this version of the LSB since there is  
59            not yet an agreed method for abstracting the implementation so that applications need  
60            not be aware of these locations during installation. However, Future Directions describes  
61            a tool, **lsbinstall**, that will make these directories implementation specific and no longer  
62            required.

## 16.2.1 File Naming Conventions

Conforming implementations and applications installing files into any of the above locations under /etc may only use filenames from the following managed namespaces:

- Assigned names. Such names must be chosen from the character set [a-z0-9]. In order to avoid conflicts these names shall be reserved through the Linux Assigned Names and Numbers Authority (LANANA). Information about the LANANA may be found at [www.lanana.org](http://www.lanana.org) (<http://www.lanana.org>).

**Note:** Commonly used names should be reserved in advance; developers for projects are encouraged to reserve names from LANANA, so that each distribution can use the same name, and to avoid conflicts with other projects.

- Hierarchical names. Script names in this category take the form:  
`<hier1>-<hier2>-...-<name>`, where name is taken from the character set [a-z0-9], and where there may be one or more `<hier-n>` components. `<hier1>` may either be an LSB provider name assigned by the LANANA, or it may be owners' DNS name in lower case, with at least one '.'. e.g. "debian.org",

78                         "staroffice.sun.com", etc. The LSB provider name assigned by LANANA shall  
79                         only consist of the ASCII characters [a-z0-9].

- 80                         • Reserved names. Names that begin with the character '\_' are reserved for  
81                         distribution use only. These names should be used for essential system packages  
82                         only.

83                         **Note:** A non-conforming application may still have polluted these managed namespaces  
84                         with unregistered filenames; a conforming application should check for namespace  
85                         collisions and take appropriate steps if they occur.

86                         In general, if a package or some system function is likely to be used on multiple systems,  
87                         the package developers or the distribution should get a registered name through  
88                         LANANA, and distributions should strive to use the same name whenever possible. For  
89                         applications which may not be essential or may not be commonly installed, the  
90                         hierarchical namespace may be more appropriate. An advantage to the hierarchical  
91                         namespace is that there is no need to consult with the LANANA before obtaining an  
92                         assigned name.

93                         Short names are highly desirable, since system administrators may need to manually  
94                         start and stop services. Given this, they should be standardized on a per-package basis.  
95                         This is the rationale behind having the LANANA organization assign these names. The  
96                         LANANA may be called upon to handle other namespace issues, such as  
97                         package/prerequisites naming.

### 16.3 User Accounting Databases

98                         The Filesystem Hierarchy Standard specifies two optional locations for user  
99                         accounting databases used by the `getutent()`, `getutent_r()`, `getutxent()`,  
100                         `getutxid()`, `getutxline()`, and `pututxline()` functions. These are  
101                         `/var/run/utmp` and `/var/run/wtmp`.

102                         The LSB does not specify the format or structure of these files, or even if they are files  
103                         at all. They should be used only as "magic cookies" to the `utmpname()` function.

### 16.4 Path For System Administration Utilities

104                         Certain utilities used for system administration (and other privileged commands)  
105                         may be stored in `/sbin`, `/usr/sbin`, and `/usr/local/sbin`. Applications requiring  
106                         to use commands identified as system administration utilities should add these  
107                         directories to their PATH. By default, as described in ISO POSIX (2003), standard  
108                         utilities shall be found on the PATH returned by `getconf PATH` (or `command -p`  
109                         `getconf PATH` to be guaranteed to invoke the correct version of `getconf`).

## 17 Additional Recommendations

### 17.1 Recommendations for applications on ownership and permissions

#### 17.1.1 Directory Write Permissions

1  
2 The application should not depend on having directory write permission in any  
directory except `/tmp`, `/var/tmp`, and the invoking user's home directory.

3 In addition, the application may store variable data in `/var/opt/package`, (where  
4 `package` is the name of the application package), if such a directory is created with  
5 appropriate permissions during the package installation.

6 For these directories the application should be able to work with directory write  
7 permissions restricted by the `S_ISVTXT` bit, implementing the restricted deletion  
8 mode as described for the XSI option for ISO POSIX (2003)..

#### 17.1.2 File Write Permissions

9 The application should not depend on file write permission to any file that it does  
10 not itself create.

#### 17.1.3 File Read and execute Permissions

11 The application should not depend on having read permission to every file and  
12 directory.

#### 17.1.4 SUID and SGID Permissions

13 The application should not depend on the set user ID or set group ID (the `S_ISUID`  
14 or `S_ISGID` permission bits) permissions of a file not packaged with the application.  
15 Instead, the distribution is responsible for assuming that all system commands have  
16 the required permissions and work correctly.

17 **Rationale:** In order to implement common security policies it is strongly advisable for  
18 applications to use the minimum set of security attributes necessary for correct operation.  
19 Applications that require substantial appropriate privilege are likely to cause problems  
20 with such security policies.

#### 17.1.5 Privileged users

21 In general, applications should not depend on running as a privileged user. This  
22 specification uses the term "appropriate privilege" throughout to identify operations  
23 that cannot be achieved without some special granting of additional privilege.

24 Applications that have a reason to run with appropriate privilege should outline this  
25 reason clearly in their documentation. Users of the application should be informed,  
26 that "this application demands security privileges, which could interfere with  
27 system security".

28 The application should not contain binary-only software that requires being run  
29 with appropriate privilege, as this makes security auditing harder or even  
30 impossible.

### 17.1.6 Changing permissions

The application shall not change permissions of files and directories that do not belong to its own package. Should an application require that certain files and directories not directly belonging to the package have a particular ownership, the application shall document this requirement, and may fail during installation if the permissions on these files is inappropriate.

### 17.1.7 Removable Media (Cdrom, Floppy, etc.)

Applications that expect to be runnable from removable media should not depend on logging in as a privileged user, and should be prepared to deal with a restrictive environment. Examples of such restrictions could be default mount options that disable set-user/group-ID attributes, disabling block or character-special files on the medium, or remapping the user and group IDs of files away from any privileged value.

**Rationale:** System vendors and local system administrators want to run applications from removable media, but want the possibility to control what the application can do.

### 17.1.8 Installable applications

Where the installation of an application needs additional privileges, it must clearly document all files and system databases that are modified outside of those in `/opt/pkg-name` and `/var/opt/pkg-name`, other than those that may be updated by system logging or auditing activities.

Without this, the local system administrator would have to blindly trust a piece of software, particularly with respect to its security.

## 18 Additional Behaviors

### 18.1 Mandatory Optional Behaviors

1 This section specifies behaviors in which there is optional behavior in one of the  
2 standards on which the LSB relies, and where the LSB requires a specific behavior.

3 **Note:** The LSB does not require the kernel to be Linux; the set of mandated options  
4 reflects current existing practice, but may be modified in future releases.

5 LSB conforming implementations shall support the following options defined  
6 within the *ISO POSIX (2003)*:

\_POSIX\_FSYNC  
\_POSIX\_MAPPED\_FILES  
\_POSIX\_MEMLOCK  
\_POSIX\_MEMLOCK\_RANGE  
\_POSIX\_MEMORY\_PROTECTION  
\_POSIX\_PRIORITY\_SCHEDULING  
\_POSIX\_REALTIME\_SIGNALS  
\_POSIX\_THREAD\_ATTR\_STACKADDR  
\_POSIX\_THREAD\_ATTR\_STACKSIZE  
\_POSIX\_THREAD\_PROCESS\_SHARED  
\_POSIX\_THREAD\_SAFE\_FUNCTIONS  
\_POSIX\_THREADS

7  
8 The `opendir()` function shall consume a file descriptor in the same fashion as  
9 `open()`, and therefore may fail with `EMFILE` or `ENFILE`.

10 The `START` and `STOP` `termios` characters shall be changeable, as described as  
11 optional behavior in the "General Terminal Interface" section of the *ISO POSIX*  
12 (*2003*).

13 The `access()` function function shall fail with `errno` set to `EINVAL` if the *amode*  
14 argument contains bits other than those set by the bitwise inclusive OR of `R_OK`, `W_OK`,  
15 `X_OK` and `F_OK`.

16 The `link()` function shall require access to the existing file in order to succeed, as  
17 described as optional behavior in the *ISO POSIX (2003)*.

18 Calling `unlink()` on a directory shall fail. Calling `link()` specifying a directory as  
19 the first argument shall fail. See also `unlink`.

20 **Note:** Linux allows `rename()` on a directory without having write access, but the LSB  
21 does not require this.

#### 18.1.1 Special Requirements

22 LSB conforming systems shall enforce certain special additional restrictions above  
23 and beyond those required by ISO POSIX (2003).

24 **Note:** These additional restrictions are required in order to support the testing and  
25 certification programs associated with the LSB. In each case, these are values that defined  
26 macros must not have; conforming applications that use these values shall trigger a  
27 failure in the interface that is otherwise described as a "may fail".

28 The `fcntl()` function shall treat the "cmd" value -1 as invalid.

29           The *whence* value -1 shall be an invalid value for the `lseek()`, `fseek()` and  
30           `fcntl()` functions.

31           The value -5 shall be an invalid signal number.

32           If the `sigaddset()` or `sigdelset()` functions are passed an invalid signal number,  
33           they shall return with `EINVAL`. Implementations are only required to enforce this  
34           requirement for signal numbers which are specified to be invalid by this  
35           specification (such as the -5 mentioned above).

36           The mode value -1 to the `access()` function shall be treated as invalid.

37           A value of -1 shall be an invalid "\_PC\_..." value for `pathconf()`.

38           A value of -1 shall be an invalid "\_SC\_..." value for `sysconf()`.

39           The *nl\_item* value -1 shall be invalid for `nl_langinfo()`.

40           The value -1 shall be an invalid "\_CS\_..." value for `confstr()`.

41           The value "a" shall be an invalid *mode* argument to `popen()`.

42           The `fcntl()` function shall fail and set `errno` to `EDEADLK` if the *cmd* argument is  
43           `F_SETLKW`, and the lock is blocked by a lock from another process already blocked by  
44           the current process.

45           The `opendir()` function shall consume a file descriptor; the `readdir()` function  
46           shall fail and set `errno` to `EBADF` if the underlying file descriptor is closed.

47           The `link()` function shall not work across file systems, and shall fail and set `errno`  
48           to `EXDEV` as described as optional behavior in ISO POSIX (2003).

## 19 Localization

### 19.1 Introduction

1 In order to install a message catalog, the installation procedure shall supply the  
2 message catalog in a format readable by the **msgfmt** utility, which shall be invoked  
3 to compile the message catalog into an appropriate binary format on the target  
4 system.

5 **Rationale:** The original intent was to allow an application to contain the binary GNU  
6 MO format files. However, the format of these files is not officially stable, hence it is  
7 necessary to compile these catalogs on the target system. These binary catalogs may  
8 differ from architecture to architecture as well.

9 The resulting binary message catalog shall be located in the package's private area  
10 under `/opt`, and the application may use `bindtextdomain()` to specify this location.

11 Implementations shall support the POSIX and C locales as specified in ISO POSIX  
12 (2003). Other locales may be supported.

13 Implementations may define additional locale categories not defined by that  
14 standard.

15 **Note:** Implementations choosing additional locale categories should be aware of  
16 ISO/IEC TR14652 and are advised not to choose names that conflict with that  
17 specification. If implementations provide locale categories whose names are part of the  
18 FDCC set of ISO/IEC TR14652, they should behave as defined by that specification.

### 19.2 Regular Expressions

19 Utilities that process regular expressions shall support Basic Regular Expressions  
20 and Extended Regular Expressions as specified in ISO POSIX (2003), with the  
21 following exceptions:

22 Range expression (such as `[a-z]`) can be based on code point order instead of  
23 collating element order.

24 Equivalence class expression (such as `[=a=]`) and multi-character collating element  
25 expression (such as `[.ch.]`) are optional.

26 Handling of a multi-character collating element is optional.

27 This affects at least the following utilities:

- 28 • **awk** (see `awk`)
- 29 • **grep** (see `grep`) (including **egrep**, see `egrep`)
- 30 • **sed** (see `sed`)

31 It also affects the behavior of interfaces in the base libraries, including at least

- 32 • `regexec()` (see `regexec`)

### 19.3 Pattern Matching Notation

33 Utilities that perform filename pattern matching (also known as Filename Globbing)  
34 shall do it as specified in ISO POSIX (2003), Pattern Matching Notation, with the  
35 following exceptions:

- 36        Pattern bracket expressions (such as [ a-z ]) can be based on code point order instead  
37        of collating element order.
- 38        Equivalence class expression (such as [=a=]) and multi-character collating element  
39        expression (such as [ .ch. ]) are optional.
- 40        Handling of a multi-character collating element is optional.
- 41        This affects at least the following utilities: **cpio** (cpio), **find** (find) and **tar** (tar).

## **VII System Initialization**

## 20 System Initialization

### 20.1 Cron Jobs

1 In addition to the individual user `crontab` files specified by ISO POSIX (2003) stored  
2 under `/var/spool/cron`, the process that executes scheduled commands shall also  
3 process the following additional `crontab` files: `/etc/crontab`, `/etc/cron.d/*`. The  
4 installation of a package shall not modify the configuration file `/etc/crontab`.

5 If a package wishes to install a job that has to be executed periodically, it shall place  
6 an executable *cron script* in one of the following directories:

7     `/etc/cron.hourly`  
8     `/etc/cron.daily`  
9     `/etc/cron.weekly`  
10    `/etc/cron.monthly`

11 As these directory names suggest, the files within them are executed on a hourly,  
12 daily, weekly, or monthly basis, respectively, under the control of an entry in one of  
13 the system `crontab` files, at an unspecified time of day. See below for the rules  
14 concerning the names of cron scripts.

15 **Note:** It is recommended that cron scripts installed in any of these directories be script  
16 files rather than compiled binaries so that they may be modified by the local system  
17 administrator. Conforming applications may only install cron scripts which use an  
18 interpreter required by this specification or provided by this or another conforming  
19 application.

20 This specification does not define the concept of a package *upgrade*. Implementations  
21 may do different things when packages are upgraded, including not replacing a cron  
22 script if it marked as a configuration file, particularly if the cron script appears to have  
23 been modified since installation. In some circumstances, the cron script may not be  
24 removed when the package is uninstalled. Applications should design their installation  
25 procedure and cron scripts to be robust in the face of such behavior. In particular, cron  
26 scripts should not fail obscurely if run in unexpected circumstances. Testing for the  
27 existence of application binaries before executing them is suggested.

28 Future versions of this specification may remove the need to install file directly into these  
29 directories, and instead abstract the interface to the `cron` utility in such a way as to hide  
30 the implementation. Please see Future Directions.

31 If a certain task has to be executed at other than the predefined frequencies, the  
32 package shall install a file `/etc/cron.d/cron-name`. The file shall have the same  
33 format as that described for the `crontab` command in ISO POSIX (2003), except that  
34 there shall be an additional field, `username`, before the name of the command to  
35 execute. For completeness, the seven fields shall be:

- 36     1. Minute [0,59]  
37     2. Hour [0,23]  
38     3. Day of the month [1,31]  
39     4. Month of the year [1,12]  
39     5. Day of the week [0,6] (with 0=Sunday)  
39     6. Username  
39     7. command [args ...]

40 This file shall be processed by the system automatically, with the named command  
 41 being run at the specified time, as the specified username.

42 Applications installing files in these directories shall use the LSB naming  
 43 conventions (see File Naming Conventions).

## 20.2 Init Script Actions

44 Conforming applications which need to execute commands on changes to the sys-  
 45 tem run level (including boot and shutdown), may install one or more *init scripts*.  
 46 Init scripts provided by conforming applications shall accept a single argument  
 47 which selects the action:

<b>start</b>	start the service
<b>stop</b>	stop the service
<b>restart</b>	stop and restart the service if the service is already running, otherwise start the service
<b>try-restart</b>	restart the service if the service is already running
<b>reload</b>	cause the configuration of the service to be reloaded without actually stopping and restarting the service
<b>force-reload</b>	cause the configuration to be reloaded if the service supports this, otherwise restart the service if it is running
<b>status</b>	print the current status of the service

48  
 49 The **start**, **stop**, **restart**, **force-reload**, and **status** actions shall be supported by all init  
 50 scripts; the **reload** and the **try-restart** actions are optional. Other init-script actions  
 51 may be defined by the init script.

52 Init scripts shall ensure that they will behave sensibly if invoked with **start** when the  
 53 service is already running, or with **stop** when not running, and that they do not kill  
 54 similarly-named user processes. The best way to achieve this is to use the init-script  
 55 functions provided by */lib/lsb/init-functions* (see Init Script Functions)

56 If a service reloads its configuration automatically (as in the case of cron, for  
 57 example), the **reload** action of the init script shall behave as if the configuration was  
 58 reloaded successfully. The **restart**, **try-restart**, **reload** and **force-reload** actions may  
 59 be atomic; that is if a service is known not to be operational after a restart or reload,  
 60 the script may return an error without any further action.

61 **Note:** This specification does not define the concept of a package *upgrade*.  
 62 Implementations may do different things when packages are upgraded, including not  
 63 replacing an init script if it is marked as a configuration file, particularly if the file  
 64 appears to have been modified since installation. In some circumstances, the init script  
 65 may not be removed when the package is uninstalled. Applications should design their  
 66 installation procedure and init scripts to be robust in the face of such behavior. In  
 67 particular, init scripts should not fail obscurely if run in unexpected circumstances.  
 68 Testing for the existence of application binaries before executing them is suggested.

69 If the **status** action is requested, the init script will return the following exit status  
 70 codes.

0	program is running or service is OK
1	program is dead and /var/run pid file exists

2	program is dead and /var/lock lock file exists
3	program is not running
4	program or service status is unknown
5-99	reserved for future LSB use
100-149	reserved for distribution use
150-199	reserved for application use
200-254	reserved
71	
72	For all other init-script actions, the init script shall return an exit status of zero if the
73	action was successful. Otherwise, the exit status shall be non-zero, as defined below.
74	In addition to straightforward success, the following situations are also to be
75	considered successful:
76	• restarting a service (instead of reloading it) with the <b>force-reload</b> argument
77	• running <b>start</b> on a service already running
78	• running <b>stop</b> on a service already stopped or not running
79	• running <b>restart</b> on a service already stopped or not running
80	• running <b>try-restart</b> on a service already stopped or not running
81	In case of an error while processing any init-script action except for <b>status</b> , the init
82	script shall print an error message and exit with a non-zero status code:
83	
1	generic or unspecified error (current practice)
2	invalid or excess argument(s)
3	unimplemented feature (for example, "reload")
4	user had insufficient privilege
5	program is not installed
6	program is not configured
7	program is not running
8-99	reserved for future LSB use
100-149	reserved for distribution use
150-199	reserved for application use
200-254	reserved
84	Error and status messages should be printed with the logging functions (see Init
85	Script Functions) <code>log_success_msg()</code> , <code>log_failure_msg()</code> and
86	<code>log_warning_msg()</code> . Scripts may write to standard error or standard output, but
87	implementations need not present text written to standard error/output to the user
88	or do anything else with it.
89	<b>Note:</b> Since init scripts may be run manually by a system administrator with
90	non-standard environment variable values for PATH, USER, LOGNAME, etc., init
91	scripts should not depend on the values of these environment variables. They should set
92	them to some known/default values if they are needed.

## 20.3 Comment Conventions for Init Scripts

Conforming applications may install one or more init scripts. These init scripts must be activated by invoking the **install\_initd** command. Prior to package removal, the changes applied by **install\_initd** must be undone by invoking **remove\_initd**. See Installation and Removal of Init Scripts for more details.

97       **install\_initd** and **remove\_initd** determine actions to take by decoding a specially  
 98        formatted block of lines in the script. This block shall be delimited by the lines

```
99        ### BEGIN INIT INFO
100      ### END INIT INFO
```

101      The delimiter lines may contain trailing whitespace, which shall be ignored. All lines  
 102        inside the block shall begin with a hash character '#' in the first column, so the shell  
 103        interprets them as comment lines which do not affect operation of the script. The  
 104        lines shall be of the form:

```
105      # {keyword}: arg1 [arg2...]
```

106      with exactly one space character between the '#' and the keyword, with a single  
 107        exception. In lines following a line containing the **Description** keyword, and until  
 108        the next keyword or block ending delimiter is seen, a line where the '#' is followed  
 109        by more than one space or a tab character shall be treated as a continuation of the  
 110        previous line.

111      The information extracted from the block is used by the installation tool or the  
 112        init-script system to assure that init scripts are run in the correct order. It is  
 113        unspecified whether the information is evaluated only when **install\_initd** runs,  
 114        when the init scripts are executed, or both. The information extracted includes run  
 115        levels, defined in Run Levels, and boot facilities, defined in Facility Names.

116      The following keywords, with their arguments, are defined:

```
117      Provides: boot_facility_1 [boot_facility_2...]
```

118      boot facilities provided by this init script. When an init script is run with a **start**  
 119        argument, the boot facility or facilities specified by the **Provides** keyword shall  
 120        be deemed present and hence init scripts which require those boot facilities  
 121        should be started later. When an init script is run with a **stop** argument, the boot  
 122        facilities specified by the **Provides** keyword are deemed no longer present.

```
123      Required-Start: boot_facility_1 [boot_facility_2...]
```

124      facilities which must be available during startup of this service. The init-script  
 125        system should insure init scripts which provide the **Required-Start** facilities are  
 126        started before starting this script.

```
127      Required-Stop: boot_facility_1 [boot_facility_2...]
```

128      facilities which must be available during the shutdown of this service. The  
 129        init-script system should avoid stopping init scripts which provide the  
 130        **Required-Stop** facilities until this script is stopped.

```
131      Should-Start: boot_facility_1 [boot_facility_2...]
```

132      facilities which, if present, should be available during startup of this service.  
 133        This allows for weak dependencies which do not cause the service to fail if a  
 134        facility is not available. The service may provide reduced functionality in this  
 135        situation. Conforming applications should not rely on the existence of this  
 136        feature.

```
137      Should-Stop: boot_facility_1 [boot_facility_2...]
```

138      facilities which should be available during shutdown of this service.

```

139      Default-Start: run_level_1 [run_level_2...]
140      Default-Stop: run_level_1 [run_level_2...]
141          which run levels should by default run the init script with a start (stop)
142          argument to start (stop) the services controlled by the init script.
143
144          For example, if a service should run in runlevels 3, 4, and 5 only, specify
145          "Default-Start: 3 4 5" and "Default-Stop: 0 1 2 6".
146
147      Short-Description: short_description
148          provide a brief description of the actions of the init script. Limited to a single
149          line of text.
150
151      Description: multiline_description
152          provide a more complete description of the actions of the init script. May span
153          multiple lines. In a multiline description, each continuation line shall begin with
154          a '#' followed by tab character or a '#' followed by at least two space characters.
155          The multiline description is terminated by the first line that does not match this
156          criteria.
157
158      Additional keywords may be defined in future versions of this specification. Also,
159      implementations may define local extensions by using the prefix X-Implementor.
160      For example, X-RedHat-foobardecl, or X-Debian-xyzzydecl.
161
162      Example:
163
164      ## BEGIN INIT INFO
165          # Provides: lsb-ourdb
166          # Required-Start: $local_fs $network $remote_fs
167          # Required-Stop: $local_fs $network $remote_fs
168          # Default-Start: 2 3 4 5
169          # Default-Stop: 0 1 6
170          # Short-Description: start and stop OurDB
171          # Description: OurDB is a very fast and reliable database
172          #                 engine used for illustrating init scripts
173          ## END INIT INFO

```

168  
169  
170  
The comment conventions described in this section are only required for init scripts  
installed by conforming applications. Conforming runtime implementations are not  
required to use this scheme in their system provided init scripts.

171  
172  
173  
**Note:** This specification does not require, but is designed to allow, the development of a  
system which runs init scripts in parallel. Hence, enforced-serialization of scripts is  
avoided unless it is explicitly necessary.

## 20.4 Installation and Removal of Init Scripts

174  
175  
176  
Conforming applications may install one or more initialization scripts (or *init scripts*).  
An init script shall be installed in `/etc/init.d` (which may be a symbolic link to  
another location), by the package installer.

177  
178  
179  
**Note:** The requirement to install scripts in `/etc/init.d` may be removed in future  
versions of this specification. See Host-specific system configuration and Future  
Directions for further details.

180  
181  
182  
During the installer's post-install processing phase the program  
`/usr/lib/lsb/install_initd` must be called to activate the init script. Activation consists  
of arranging for the init script to be called in the correct order on system run-level

183 changes (including system boot and shutdown), based on dependencies supplied in  
 184 the init script (see Comment Conventions for Init Scripts). The **install\_initd**  
 185 command should be thought of as a wrapper which hides the implementation  
 186 details; how any given implementation arranges for the init script to be called at the  
 187 appropriate time is not specified.

188 Example: if an init script specified "Default-Start: 3 4 5" and "Default-Stop: 0 1 2 6",  
 189 **install\_initd** might create "start" symbolic links with names starting with 'S' in  
 190 /etc/rc3.d, /etc/rc4.d and /etc/rc5.d and "stop" symbolic links with names starting  
 191 with 'K' in /etc/rc0.d, /etc/rc1.d, /etc/rc2.d and /etc/rc6.d. Such a scheme would  
 192 be similar to the System V Init mechanism, but is by no means the only way this  
 193 specification could be implemented.

194 The **install\_initd** command takes a single argument, the full pathname of the  
 195 installed init script. The init script must already be installed in /etc/init.d. The  
 196 **install\_initd** command will not copy it there, only activate it once it has been  
 197 installed. For example:

```
198   /usr/lib/lsb/install_initd /etc/init.d/example.com-coffeed
```

199 The **install\_initd** command shall return an exit status of zero if the init-script  
 200 activation was successful or if the init script was already activated. If the  
 201 dependencies in the init script (see Comment Conventions for Init Scripts) cannot be  
 202 met, an exit status of one shall be returned and the init script shall not be activated.

203 When a software package is removed, **/usr/lib/lsb/remove\_initd** must be called to  
 204 deactivate the init script. This must occur before the init script itself is removed, as  
 205 the dependency information in the script may be required for successful completion.  
 206 Thus the installer's pre-remove processing phase must call **remove\_initd**, and pass  
 207 the full pathname of the installed init script. The package installer is still responsible  
 208 for removing the init script. For example:

```
209   /usr/lib/lsb/remove_initd /etc/init.d/example.com-coffeed
```

210 The **remove\_initd** program shall return an exit status of zero if the init script has  
 211 been successfully deactivated or if the init script is not activated. If another init script  
 212 which depends on a boot facility provided by this init script is activated, an exit  
 213 status of one shall be returned and the init script shall remain activated. The installer  
 214 must fail on such an exit code so it does not subsequently remove the init script.

215 **Note:** This specification does not describe a mechanism for the system administrator to  
 216 manipulate the run levels at which an init script is started or stopped. There is no  
 217 assurance that modifying the comment block for this purpose will have the desired  
 218 effect.

## 20.5 Run Levels

219 The following *run levels* are specified for use by the **Default-Start** and **Default-Stop**  
 220 actions defined in Comment Conventions for Init Scripts as hints to the **install\_initd**  
 221 command. Conforming implementations are not required to provide these exact run  
 222 levels or give them the meanings described here, and may map any level described  
 223 here to a different level which provides the equivalent functionality. Applications  
 224 may not depend on specific run-level numbers.

0	halt
1	single user mode
2	multiuser with no network services

		exported
3		normal/full multiuser
4		reserved for local use, default is normal/full multiuser
5		multiuser with a display manager equivalent
6		reboot

226           **Note:** These run levels were chosen as reflecting the most frequent existing practice, and  
227            in the absence of other considerations, implementors are strongly encouraged to follow  
228            this convention to provide consistency for system administrators who need to work with  
229            multiple distributions.

## 20.6 Facility Names

230 Boot facilities are used to indicate dependencies in initialization scripts, as defined in  
231 Comment Conventions for Init Scripts. Facility names are assigned to scripts by the  
232 **Provides:** keyword. Facility names that begin with a dollar sign ('\$') are reserved  
233 system facility names.

**Note:** Facility names are only recognized in the context of the init script comment block and are not available in the body of the init script. In particular, the use of the leading '\$' character does not imply system facility names are subject to shell variable expansion, since they appear inside comments.

238 Conforming applications shall not provide facilities that begin with a dollar sign.  
239 Implementations shall provide the following facility names:

240 \$local\_fs

241 all local file systems are mounted

242 \$network

basic networking support is available. Example: a server program could listen on a socket.

245 \$named

IP name-to-address translation, using the interfaces described in this specification, are available to the level the system normally provides them. Example: if a DNS query daemon normally provides this facility, then that daemon has been started.

## 250 \$portmap

daemons providing SunRPC/ONCRPC portmapping service as defined in RFC 1833: Binding Protocols for ONC RPC Version 2 (if present) are running.

253                   \$remote\_fs

254 all remote file systems are available. In some configurations, file systems such as  
255 /usr may be remote. Many applications that require **\$local\_fs** will probably  
256 also require **\$remote\_fs**.

257 \$syslog

258 system logger is operational.

259           \$**time**

260           the system time has been set, for example by using a network-based time  
 261           program such as **ntp** or **rdate**, or via the hardware Real Time Clock.

262           Other (non-system) facilities may be defined by other conforming applications.  
 263           These facilities shall be named using the same conventions defined for naming init  
 264           scripts (see Script Names). Commonly, the facility provided by a conforming init  
 265           script will have the same name as the name assigned to the init script.

## 20.7 Script Names

266           Since init scripts live in a single directory, they must share a single namespace. To  
 267           avoid conflicts, applications installing files in this directory shall use the LSB  
 268           naming conventions (see File Naming Conventions).

## 20.8 Init Script Functions

269           Each conforming init script shall execute the commands in the file  
 270           **/lib/lsb/init-functions** in the current environment (see shell special built-in  
 271           command **dot**). This file shall cause the following shell script commands to be  
 272           defined in an unspecified manner.

273           **Note:** This can be done either by adding a directory to the PATH variable which defines  
 274           these commands, or by defining shell aliases or functions.

275           Although the commands made available via this mechanism need not be conforming  
 276           applications in their own right, applications that use them should only depend on  
 277           features described in this specification.

278           Conforming scripts shall not specify the "exit on error" option (i.e. **set -e**) when  
 279           sourcing this file, or calling any of the commands thus made available.

280           The **start\_daemon**, **killproc** and **pidofproc** functions shall use the following  
 281           algorithm for determining the status and the process identifiers of the specified  
 282           program.

- 283           1. If the **-p pidfile** option is specified, and the named **pidfile** exists, a single  
              line at the start of the **pidfile** shall be read. If this line contains one or more  
              numeric values, separated by spaces, these values shall be used. If the **-p**  
              **pidfile** option is specified and the named **pidfile** does not exist, the  
              functions shall assume that the daemon is not running.
- 288           2. Otherwise, **/var/run/basename.pid** shall be read in a similar fashion. If this  
              contains one or more numeric values on the first line, these values shall be used.  
              Optionally, implementations may use unspecified additional methods to locate  
              the process identifiers required.

292           The method used to determine the status is implementation defined, but should  
 293           allow for non-binary programs.

294           **Note:** Commonly used methods check either for the existence of the **/proc/pid** directory  
 295           or use **/proc/pid/exe** and **/proc/pid/cmdline**. Relying only on **/proc/pid/exe** is  
 296           discouraged since this specification does not specify the existence of, or semantics for,  
 297           **/proc**. Additionally, using **/proc/pid/exe** may result in a not-running status for  
 298           daemons that are written in a script language.

299           Conforming implementations may use other mechanisms besides those based on  
 300           **pidfiles**, unless the **-p pidfile** option has been used. Conforming applications  
 301           should not rely on such mechanisms and should always use a **pidfile**. When a

302           program is stopped, it should delete its `pidfile`. Multiple process identifiers shall  
 303           be separated by a single space in the `pidfile` and in the output of `pidofproc`.

304       **`start_daemon`** [`-f`] [`-n nicelevel`] [`-p pidfile`] `pathname` [`args...`]  
 305           runs the specified program as a daemon. The `start_daemon` function shall check  
 306           if the program is already running using the algorithm given above. If so, it shall  
 307           not start another copy of the daemon unless the `-f` option is given. The `-n`  
 308           option specifies a nice level. See `nice`. `start_daemon` shall return the LSB defined  
 309           exit status codes. It shall return 0 if the program has been successfully started or  
 310           is running and not 0 otherwise.

311       **`killproc`** [`-p pidfile`] `pathname` [`signal`]  
 312           The `killproc` function shall stop the specified program. The program is found  
 313           using the algorithm given above. If a signal is specified, using the `-signal_name`  
 314           or `-signal_number` syntaxes as specified by the `kill` command, the program is  
 315           sent that signal. Otherwise, a `SIGTERM` followed by a `SIGKILL` after an  
 316           unspecified number of seconds shall be sent. If a program has been terminated,  
 317           the `pidfile` should be removed if the terminated process has not already done  
 318           so. The `killproc` function shall return the LSB defined exit status codes. If called  
 319           without a signal, it shall return 0 if the program has been stopped or is not  
 320           running and not 0 otherwise. If a signal is given, it shall return 0 only if the  
 321           program is running.

322       **`pidofproc`** [`-p pidfile`] `pathname`  
 323           The `pidofproc` function shall return one or more process identifiers for a  
 324           particular daemon using the algorithm given above. Only process identifiers of  
 325           running processes should be returned. Multiple process identifiers shall be  
 326           separated by a single space.

327       **Note:** A process may exit between `pidofproc` discovering its identity and the caller of  
 328           `pidofproc` being able to act on that identity. As a result, no test assertion can be  
 329           made that the process identifiers returned by `pidofproc` shall be running processes.

330       The `pidofproc` function shall return the LSB defined exit status codes for  
 331           "status". It shall return 0 if the program is running and not 0 otherwise.

332       **`log_success_msg`** `message`  
 333           The `log_success_msg` function shall cause the system to write a success message  
 334           to an unspecified log file. The format of the message is unspecified. The  
 335           `log_success_msg` function may also write a message to the standard output.

336       **Note:** The message should be relatively short; no more than 60 characters is highly  
 337           desirable.

338       **`log_failure_msg`** `message`  
 339           The `log_failure_msg` function shall cause the system to write a failure message  
 340           to an unspecified log file. The format of the message is unspecified. The  
 341           `log_failure_msg` function may also write a message to the standard output.

342       **Note:** The message should be relatively short; no more than 60 characters is highly  
 343           desirable.

344           **log\_warning\_msg** message

345           The **log\_warning\_msg** function shall cause the system to write a warning  
346           message to an unspecified log file. The format of the message is unspecified.  
347           The **log\_warning\_msg** function may also write a message to the standard  
348           output.

349           **Note:** The message should be relatively short; no more than 60 characters is highly  
350           desirable.

## **VIII Users & Groups**

## 21 Users & Groups

### 21.1 User and Group Database

1 The format of the User and Group databases is not specified. Programs may only  
2 read these databases using the provided API. Changes to these databases should be  
3 made using the provided commands.

### 21.2 User & Group Names

4 Table 21-1 describes required mnemonic user and group names. This specification  
5 makes no attempt to numerically assign user or group identity numbers, with the  
6 exception that both the User ID and Group ID for the user `root` shall be equal to 0.

7 **Table 21-1 Required User & Group Names**

User	Group	Comments
root	root	Administrative user with all appropriate privileges
bin	bin	Legacy User ID/Group ID <sup>a</sup>
daemon	daemon	Legacy User ID/Group ID <sup>b</sup>

8 Notes:

- a The `bin` User ID/Group ID is included for compatibility with legacy applications. New applications should no longer use the `bin` User ID/Group ID.
- b The `daemon` User ID/Group ID was used as an unprivileged User ID/Group ID for daemons to execute under in order to limit their access to the system. Generally daemons should now run under individual User ID/Group IDs in order to further partition daemons from one another.

9 Table 21-2 is a table of optional mnemonic user and group names. This specification  
10 makes no attempt to numerically assign uid or gid numbers. If the username exists  
11 on a system, then they should be in the suggested corresponding group. These user  
12 and group names are for use by distributions, not by applications.

13 **Table 21-2 Optional User & Group Names**

User	Group	Comments
adm	adm	Administrative special privileges
lp	lp	Printer special privileges
sync	sync	Login to sync the system
shutdown	shutdown	Login to shutdown the system
halt	halt	Login to halt the system

User	Group	Comments
mail	mail	Mail special privileges
news	news	News special privileges
uucp	uucp	UUCP special privileges
operator	root	Operator special privileges
man	man	Man special privileges
nobody	nobody	Used by NFS

14

15 Only a minimum working set of "user names" and their corresponding "user groups"  
 16 are required. Applications cannot assume non system user or group names will be  
 17 defined.

18 Applications cannot assume any policy for the default file creation mask (**umask**) or  
 19 the default directory permissions a user may have. Applications should enforce user  
 20 only file permissions on private files such as mailboxes. The location of the users  
 21 home directory is also not defined by policy other than the recommendations of the  
 22 Filesystem Hierarchy Standard and should be obtained by the `getpwnam()`,  
 23 `getpwnam_r()`, `getpwent()`, `getpwuid()`, and `getpwuid_r()` functions.

## 21.3 User ID Ranges

24 The system User IDs from 0 to 99 should be statically allocated by the system, and  
 25 shall not be created by applications.

26 The system User IDs from 100 to 499 should be reserved for dynamic allocation by  
 27 system administrators and post install scripts using **useradd**.

## 21.4 Rationale

28 The purpose of specifying optional users and groups is to reduce the potential for  
 29 name conflicts between applications and distributions.

## **IX Package Format and Installation**

## 22 Software Installation

### 22.1 Introduction

1 Applications shall either be packaged in the RPM packaging format as defined in  
2 this specification, or supply an installer which is LSB conforming (for example, calls  
3 LSB commands and utilities).

4 **Note:** Supplying an RPM format package is encouraged because it makes systems easier  
5 to manage. This specification does not require the implementation to use RPM as the  
6 package manager; it only specifies the format of the package file.

7 Applications are also encouraged to uninstall cleanly.

8 A package in RPM format may include a dependency on the LSB Core and other LSB  
9 specifications, as described in Section 22.6. Packages that are not in RPM format may  
10 test for the presence of a conforming implementation by means of the **lsb\_release**  
11 utility.

12 Implementations shall provide a mechanism for installing applications in this  
13 packaging format with some restrictions listed below.

14 **Note:** The implementation itself may use a different packaging format for its own  
15 packages, and of course it may use any available mechanism for installing the  
16 LSB-conformant packages.

### 22.2 Package File Format

17 An RPM format file consists of 4 sections, the Lead, Signature, Header, and the  
18 Payload. All values are stored in network byte order.

19 **Table 22-1 RPM File Format**

Lead
Signature
Header
Payload

21 These 4 sections shall exist in the order specified.

22 The lead section is used to identify the package file.

23 The signature section is used to verify the integrity, and optionally, the authenticity  
24 of the majority of the package file.

25 The header section contains all available information about the package. Entries  
26 such as the package's name, version, and file list, are contained in the header.

27 The payload section holds the files to be installed.

#### 22.2.1 Lead Section

```
28 struct rpmlead {  
29     unsigned char magic[4];  
30     unsigned char major, minor;  
31     short type;  
32     short archnum;
```

```

33         char name[66];
34         short osnum;
35         short signature_type;
36         char reserved[16];
37     } ;
38
39     magic
40         Value identifying this file as an RPM format file. This value shall be
41             "\355\253\356\333".
42
43     major
44         Value indicating the major version number of the file format version. This value
45             shall be 3.
46
47     minor
48         Value indicating the minor revision number of file format version. This value
49             shall be 0.
50
51     type
52         Value indicating whether this is a source or binary package. This value shall be
53             0 to indicate a binary package.
54
55     archnum
56         Value indicating the architecture for which this package is valid. This value is
57             specified in the architecture specific supplement.
58
59     name
60         A NUL terminated string that provides the package name. This name shall
61             conform with the Package Naming section of this specification.
62
63     osnum
64         Value indicating the Operating System for which this package is valid. This
65             value shall be 1.
66
67     signature_type
68         Value indicating the type of the signature used in the Signature part of the file.
69             This value shall be 5.
70
71     reserved
72         Reserved space. The value is undefined.

```

## 22.2.2 Header Structure

The Header structure is used for both the Signature and Header Sections. A Header Structure consists of 3 parts, a Header record, followed by 1 or more Index records, followed by 0 or more bytes of data associated with the Index records. A Header structure shall be aligned to an 8 byte boundary.

**Table 22-2 Signature Format**

Header Record
Array of Index Records

## Store of Index Values

## **22.2.2.1 Header Record**

```
71         struct rpmheader {
72             unsigned char magic[4];
73             unsigned char reserved[4];
74             int nindex;
75             int hsize;
76             } ;
```

77 *magic*

78                   Value identifying this record as an RPM header record. This value shall be  
79                   "\216\255\350\001".

80 *reserved*

81                    Reserved space. This value shall be "\000\000\000\000".

82 *nindex*

83 The number of Index Records that follow this Header Record. There should be  
84 at least 1 Index Record.

85                    *hsiz*

86 The size in bytes of the storage area for the data pointed to by the Index  
87 Records.

## **22.2.2.2 Index Record**

```
89         struct rpmhdrindex {  
90             int tag;  
91             int type;  
92             int offset;  
93             int count;  
94         } ;
```

95 tag

96 Value identifying the purpose of the data associated with this Index Record. The  
97 value of this field is dependent on the context in which the Index Record is used,  
98 and is defined below and in later sections.

99                      *type*

Value identifying the type of the data associated with this Index Record. The possible *type* values are defined below.

102 *offset*

103 Location in the Store of the data associated with this Index Record. This value  
104 should between 0 and the value contained in the *hsize* of the Header Structure.

105 *count*

106 Size of the data associated with this Index Record. The *count* is the number of  
107 elements whose size is defined by the type of this Record.

#### **22.2.2.2.1 Index Type Values**

109 The possible values for the *type* field are defined in this table.

110

**Table 22-3 Index Type values**

Type	Value	Size (in bytes)	Alignment
RPM_NULL_TYPE	0	Not Implemented.	
RPM_CHAR_TYPE	1	1	1
RPM_INT8_TYPE	2	1	1
RPM_INT16_TYPE	3	2	2
RPM_INT32_TYPE	4	4	4
RPM_INT64_TYPE	5	Reserved.	
RPM_STRING_TYPE	6	variable, NUL terminated	1
RPM_BIN_TYPE	7	1	1
RPM_STRING_ARRAY_TYPE	8	Variable, sequence of NUL terminated strings	1
RPM_I18NSTRING_TYPE	9	variable, sequence of NUL terminated strings	1

111

The string arrays specified for entries of type RPM\_STRING\_ARRAY\_TYPE and RPM\_I18NSTRING\_TYPE are vectors of strings in a contiguous block of memory, each element separated from its neighbors by a NUL character.

115  
116  
117

Index records with type RPM\_I18NSTRING\_TYPE shall always have a *count* of 1. The array entries in an index of type RPM\_I18NSTRING\_TYPE correspond to the locale names contained in the RPMTAG\_HDRI18NTABLE index.

118  
119  
120

#### 22.2.2.2 Index Tag Values

Some values are designated as header private, and may appear in any header structure. These are defined here. Additional values are defined in later sections.

121

**Table 22-4 Header Private Tag Values**

Name	Tag Value	Type	Count	Status
RPMTAG_HEADERSIGNATURES	62	BIN	16	Optional
RPMTAG_HEADERIMMUTABLE	63	BIN	16	Optional
RPMTAG_HEADERI18NTABLE	100	STRING_ARRAY		Optional

122

#### RPMTAG\_HEADERSIGNATURES

123  
124  
125

The signature tag differentiates a signature header from a metadata header, and identifies the original contents of the signature header.

126            RPMTAG\_HEADERIMMUTABLE  
 127            This tag contains an index record which specifies the portion of the Header  
 128            Record which was used for the calculation of a signature. This data shall be  
 129            preserved or any header-only signature will be invalidated.

130            RPMTAG\_HEADERI18NTABLE  
 131            Contains a list of locales for which strings are provided in other parts of the  
 132            package.  
 133            Not all Index records defined here will be present in all packages. Each tag value has  
 134            a status which is defined here.

135            Required  
 136            This Index Record shall be present.

137            Optional  
 138            This Index Record may be present.

139            Informational  
 140            This Index Record may be present, but does not contribute to the processing of  
 141            the package.

142            Deprecated  
 143            This Index Record should not be present.

144            Obsolete  
 145            This Index Record shall not be present.

146            Reserved  
 147            This Index Record shall not be present.

### 22.2.2.3 Header Store

149            The header store contains the values specified by the Index structures. These values  
 150            are aligned according to their type and padding is used if needed. The store is  
 151            located immediately following the Index structures.

### 22.2.3 Signature Section

152            The Signature section is implemented using the Header structure. The signature  
 153            section defines the following additional tag values which may be used in the Index  
 154            structures.

155            These values exist to provide additional information about the rest of the package.

156            **Table 22-5 Signature Tag Values**

Name	Tag Value	Type	Count	Status
RPMSIGTAG_SIZE	1000	INT32	1	Required
RPMSIGTAG_PAYLOADSIZE	1007	INT32	1	Optional

158 RPMSIGTAG\_SIZE  
 159       This tag specifies the combined size of the Header and Payload sections.  
 160 RPMSIGTAG\_PAYLOADSIZE  
 161       This tag specifies the uncompressed size of the Payload archive, including the  
 162        cpio headers.  
 163 These values exist to ensure the integrity of the rest of the package.

164 **Table 22-6 Signature Digest Tag Values**

Name	Tag Value	Type	Count	Status
RPMSIGTAG_SHA1	269	STRING	1	Optional
RPMSIGTAG_MD5	1004	BIN	16	Required

165  
 166 RPMSIGTAG\_SHA1  
 167       This index contains the SHA1 checksum of the entire Header Section, including  
 168        the Header Record, Index Records and Header store.  
 169 RPMSIGTAG\_MD5  
 170       This tag specifies the 128-bit MD5 checksum of the combined Header and  
 171        Archive sections.  
 172 These values exist to provide authentication of the package.

173 **Table 22-7 Signature Signing Tag Values**

Name	Tag Value	Type	Count	Status
RPMSIGTAG_DSA	267	BIN	1	Optional
RPMSIGTAG_RSA	268	BIN	1	Optional
RPMSIGTAG_PGP	1002	BIN	1	Optional
RPMSIGTAG_GPG	1005	BIN	65	Optional

174  
 175 RPMSIGTAG\_DSA  
 176       The tag contains the DSA signature of the Header section. The data is formatted  
 177        as a Version 3 Signature Packet as specified in RFC 2440: OpenPGP Message  
 178        Format. If this tag is present, then the SIGTAG\_GPG tag shall also be present.  
 179 RPMSIGTAG\_RSA  
 180       The tag contains the RSA signature of the Header section. The data is formatted  
 181        as a Version 3 Signature Packet as specified in RFC 2440: OpenPGP Message  
 182        Format. If this tag is present, then the SIGTAG\_PGP shall also be present.

183 RPMSIGTAG\_PGP  
 184 This tag specifies the RSA signature of the combined Header and Payload  
 185 sections. The data is formatted as a Version 3 Signature Packet as specified in  
 186 RFC 2440: OpenPGP Message Format.  
 187 RPMSIGTAG\_GPG  
 188 The tag contains the DSA signature of the combined Header and Payload  
 189 sections. The data is formatted as a Version 3 Signature Packet as specified in  
 190 RFC 2440: OpenPGP Message Format.

## 22.2.4 Header Section

191 The Header section is implemented using the Header structure. The Header section  
 192 defines the following additional tag values which may be used in the Index  
 193 structures.

### 22.2.4.1 Package Information

195 The following tag values are used to indicate information that describes the package  
 196 as a whole.

197 **Table 22-8 Package Info Tag Values**

Name	Tag Value	Type	Count	Status
RPMTAG_NAME	1000	STRING	1	Required
RPMTAG_VERSION	1001	STRING	1	Required
RPMTAG_RELEASE	1002	STRING	1	Required
RPMTAG_SUMMARY	1004	I18NSTRING	1	Required
RPMTAG_DESCRIPTION	1005	I18NSTRING	1	Required
RPMTAG_SIZE	1009	INT32	1	Required
RPMTAG_DISTRIBUTION	1010	STRING	1	Informational
RPMTAG_VENDOR	1011	STRING	1	Informational
RPMTAG_LICENSE	1014	STRING	1	Required
RPMTAG_PACKAGER	1015	STRING	1	Informational
RPMTAG_GROUP	1016	I18NSTRING	1	Required
RPMTAG_URL	1020	STRING	1	Informational
RPMTAG_OS	1021	STRING	1	Required
RPMTAG_ARCH	1022	STRING	1	Required

Name	Tag Value	Type	Count	Status
RPMTAG_SOURCE	1044	STRING	1	Informational
RPMTAG_ARCHIVESIZE	1046	INT32	1	Optional
RPMTAG_RPMVERS	1064	STRING	1	Informational
RPMTAG_COOKIE	1094	STRING	1	Optional
RPMTAG_DISTURL	1123	STRING	1	Informational
RPMTAG_PAYLOADFORMAT	1124	STRING	1	Required
RPMTAG_PAYLOADCOMPRESSOR	1125	STRING	1	Required
RPMTAG_PAYLOADFLAGS	1126	STRING	1	Required

198

199 RPMTAG\_NAME

200 This tag specifies the name of the package.

201 RPMTAG\_VERSION

202 This tag specifies the version of the package.

203 RPMTAG\_RELEASE

204 This tag specifies the release of the package.

205 RPMTAG\_SUMMARY

206 This tag specifies the summary description of the package. The summary value  
207 pointed to by this index record contains a one line description of the package.

208 RPMTAG\_DESCRIPTION

209 This tag specifies the description of the package. The description value pointed  
210 to by this index record contains a full description of the package.

211 RPMTAG\_SIZE

212 This tag specifies the sum of the sizes of the regular files in the archive.

213 RPMTAG\_DISTRIBUTION

214 A string containing the name of the distribution on which the package was  
215 built.

216 RPMTAG\_VENDOR

217 A string containing the name of the organization that produced the package.

218 RPMTAG\_LICENSE

219 This tag specifies the license which applies to this package.

```

220      RPMTAG_PACKAGER
221          A string identifying the tool used to build the package.

222      RPMTAG_GROUP
223          This tag specifies the administrative group to which this package belongs.

224      RPMTAG_URL
225          Generic package information URL

226      RPMTAG_OS
227          This tag specifies the OS of the package. The OS value pointed to by this index
228          record shall be "linux".

229      RPMTAG_ARCH
230          This tag specifies the architecture of the package. The architecture value pointed
231          to by this index record is defined in architecture specific LSB specification.

232      RPMTAG_SOURCERPM
233          This tag specifies the name of the source RPM

234      RPMTAG_ARCHIVESIZE
235          This tag specifies the uncompressed size of the Payload archive, including the
236          cpio headers.

237      RPMTAG_RPMVERSION
238          This tag indicates the version of RPM tool used to build this package. The value
239          is unused.

240      RPMTAG_COOKIE
241          This tag contains an opaque string whose contents are undefined.

242      RPMTAG_DISTURL
243          URL for package

244      RPMTAG_PAYLOADFORMAT
245          This tag specifies the format of the Archive section. The format value pointed to
246          by this index record shall be 'cpio'.

247      RPMTAG_PAYLOADCOMPRESSOR
248          This tag specifies the compression used on the Archive section. The
249          compression value pointed to by this index record shall be 'gzip'

250      RPMTAG_PAYLOADFLAGS
251          This tag indicates the compression level used for the Payload. This value shall
252          always be '9'.
```

**22.2.4.2 Installation Information**

The following tag values are used to provide information needed during the installation of the package.

256

**Table 22-9 Installation Tag Values**

Name	Tag Value	Type	Count	Status
RPMTAG_PREIN	1023	STRING	1	Optional
RPMTAG_POSTIN	1024	STRING	1	Optional
RPMTAG_PREUN	1025	STRING	1	Optional
RPMTAG_POSTUN	1026	STRING	1	Optional
RPMTAG_PREINPROG	1085	STRING	1	Optional
RPMTAG_POSTINPROG	1086	STRING	1	Optional
RPMTAG_PREUNPROG	1087	STRING	1	Optional
RPMTAG_POSTUNPROG	1088	STRING	1	Optional

257

258 RPMTAG\_PREIN

259 This tag specifies the preinstall scriptlet. If present, then  
 260 RPMTAG\_PREINPROG shall also be present.

261

262 RPMTAG\_POSTIN

263 This tag specifies the postinstall scriptlet. If present, then  
 RPMTAG\_POSTINPROG shall also be present.

264

265 RPMTAG\_PREUN

266 This tag specifies the preuninstall scriptlet. If present, then  
 RPMTAG\_PREUNPROG shall also be present.

267

268 RPMTAG\_POSTUN

269 This tag specified the postuninstall scriptlet. If present, then  
 RPMTAG\_POSTUNPROG shall also be present.

270

271 RPMTAG\_PREINPROG

272 This tag specifies the name of the interpreter to which the preinstall scriptlet will  
 be passed. The interpreter pointed to by this index record shall be /bin/sh.

273

274 RPMTAG\_POSTINPROG

275 This tag specifies the name of the interpreter to which the postinstall scriptlet  
 will be passed. The interpreter pointed to by this index record shall be /bin/sh.

276

277 RPMTAG\_PREUNPROG

278 This tag specifies the name of the interpreter to which the preuninstall scriptlet  
 will be passed. The interpreter pointed to by this index record shall be /bin/sh.

279                   RPMTAG\_POSTUNPROG  
 280                   This program specifies the name of the interpreter to which the postuninstall  
 281                   scriptlet will be passed. The interpreter pointed to by this index record shall be  
 282                   /bin/sh.

### 22.2.4.3 File Information

284                   The following tag values are used to provide information about the files in the  
 285                   payload. This information is provided in the header to allow more efficient access of  
 286                   the information.

287                   **Table 22-10 File Info Tag Values**

Name	Tag Value	Type	Count	Status
RPMTAG_OLDFILENAMES	1027	STRING_ARRAY		Optional
RPMTAG_FILESIZES	1028	INT32		Required
RPMTAG_FILEMODES	1030	INT16		Required
RPMTAG_FILERDEVS	1033	INT16		Required
RPMTAG_FILEMTIMES	1034	INT32		Required
RPMTAG_FILEMD5S	1035	STRING_ARRAY		Required
RPMTAG_FILELINKTOS	1036	STRING_ARRAY		Required
RPMTAG_FILEFLAGS	1037	INT32		Required
RPMTAG_FILEUSERNAME	1039	STRING_ARRAY		Required
RPMTAG_FILEGROUPNAME	1040	STRING_ARRAY		Required
RPMTAG_FILEDVICES	1095	INT32		Required
RPMTAG_FILEINODES	1096	INT32		Required
RPMTAG_FILELANGS	1097	STRING_ARRAY		Required
RPMTAG_DIRINDEXES	1116	INT32		Optional
RPMTAG_BASENAMES	1117	STRING_ARRAY		Optional
RPMTAG_DIRNAMES	1118	STRING_ARRAY		Optional

289 RPMTAG\_OLDFILENAMES  
 290     This tag specifies the filenames when not in a compressed format as determined  
 291     by the absence of rpmlib(CompressedFileNames) in the  
 292     RPMTAG\_REQUIRENAME index.

293 RPMTAG\_FILESIZES  
 294     This tag specifies the size of each file in the archive.

295 RPMTAG\_FILEMODES  
 296     This tag specifies the mode of each file in the archive.

297 RPMTAG\_FILERDEVS  
 298     This tag specifies the device number from which the file was copied.

299 RPMTAG\_FILEMTIMES  
 300     This tag specifies the modification time in seconds since the epoch of each file in  
 301     the archive.

302 RPMTAG\_FILEMD5S  
 303     This tag specifies the ASCII representation of the MD5 sum of the  
 304     corresponding file contents. This value is empty if the corresponding archive  
 305     entry is not a regular file.

306 RPMTAG\_FILELINKTOS  
 307     The target for a symlink, otherwise NULL.

308 RPMTAG\_FILEFLAGS  
 309     This tag specifies the bit(s) to classify and control how files are to be installed.  
 310     See below.

311 RPMTAG\_FILEUSERNAME  
 312     This tag specifies the owner of the corresponding file.

313 RPMTAG\_FILEGROUPNAME  
 314     This tag specifies the group of the corresponding file.

315 RPMTAG\_FILEDEVICES  
 316     This tag specifies the 16 bit device number from which the file was copied.

317 RPMTAG\_FILEINODES  
 318     This tag specifies the inode value from the original file system on the system  
 319     on which it was built.

320 RPMTAG\_FILELANGS  
 321     This tag specifies a per-file locale marker used to install only locale specific  
 322     subsets of files when the package is installed.

323 RPMTAG\_DIRINDEXES  
 324     This tag specifies the index into the array provided by the  
 325     RPMTAG\_DIRNAMES Index which contains the directory name for the  
 326     corresponding filename.

327 RPMTAG\_BASENAMES  
 328       This tag specifies the base portion of the corresponding filename.  
 329 RPMTAG\_DIRNAMES  
 330  
 331 One of RPMTAG\_OLDFILENAMES or the tuple  
 332 RPMTAG\_DIRINDEXES, RPMTAG\_BASENAMES, RPMTAG\_DIRNAMES shall be present, but  
 333 not both.

334 **22.2.4.3.1 File Flags**

335 The RPMTAG\_FILEFLAGS tag value shall identify various characteristics of the file in  
 336 the payload that it describes. It shall be an INT32 value consisting of either the value  
 337 RPMFILE\_NONE (0) or the bitwise inclusive or of one or more of the following values:

338 **Table 22-11 File Flags**

Name	Value
RPMFILE_CONFIG	(1 << 0)
RPMFILE_DOC	(1 << 1)
RPMFILE_DONOTUSE	(1 << 2)
RPMFILE_MISSINGOK	(1 << 3)
RPMFILE_NOREPLACE	(1 << 4)
RPMFILE_SPECFILE	(1 << 5)
RPMFILE_GHOST	(1 << 6)
RPMFILE_LICENSE	(1 << 7)
RPMFILE_README	(1 << 8)
RPMFILE_EXCLUDE	(1 << 9)

339

340 These bits have the following meaning:

341 RPMFILE\_CONFIG

342       The file is a configuration file, and an existing file should be saved during a  
 343 package upgrade operation and not removed during a package removal  
 344 operation.

345 RPMFILE\_DOC

346       The file contains documentation.

347 RPMFILE\_DONOTUSE

348       This value is reserved for future use; conforming packages may not use this  
 349 flag.

350 RPMFILE\_MISSINGOK

351       The file need not exist on the installed system.

352 RPMFILE\_NOREPLACE  
 353     Similar to the RPMFILE\_CONFIG, this flag indicates that during an upgrade  
 354       operation the original file on the system should not be altered.

355 RPMFILE\_SPECFILE  
 356     The file is a package specification.

357 RPMFILE\_GHOST  
 358     The file is not actually included in the payload, but should still be considered as  
 359       a part of the package. For example, a log file generated by the application at run  
 360       time.

361 RPMFILE\_LICENSE  
 362     The file contains the license conditions.

363 RPMFILE\_README  
 364     The file contains high level notes about the package.

365 RPMFILE\_EXCLUDE  
 366     The corresponding file is not a part of the package, and should not be installed.

#### 22.2.4.4 Dependency Information

The following tag values are used to provide information about interdependencies between packages.

**Table 22-12 Package Dependency Tag Values**

Name	Tag Value	Type	Count	Status
RPMTAG_PROV_IDENAME	1047	STRING_ARRAY	1	Required
RPMTAG_REQUIREFLAGS	1048	INT32		Required
RPMTAG_REQUIRENAME	1049	STRING_ARRAY		Required
RPMTAG_REQUIREVERSION	1050	STRING_ARRAY		Required
RPMTAG_CONF_LICTFLAGS	1053	INT32		Optional
RPMTAG_CONF_LICTNAME	1054	STRING_ARRAY		Optional
RPMTAG_CONF_LICTVERSION	1055	STRING_ARRAY		Optional
RPMTAG_OBSOLETENAME	1090	STRING_ARRAY		Optional
RPMTAG_PROVIDEFLAGS	1112	INT32		Required
RPMTAG_PROVIDE	1113	STRING_ARRAY		Required

Name	Tag Value	Type	Count	Status
IDEVERSION		RAY		
RPMTAG_OBSOLETEFLAGS	1114	INT32	1	Optional
RPMTAG_OBSOLETEREVISION	1115	STRING_AR RAY		Optional

371

372 RPMTAG\_PROVIDENAME

373

This tag indicates the name of the dependency provided by this package.

374

RPMTAG\_REQUIREFLAGS

375

Bits(s) to specify the dependency range and context.

376

RPMTAG\_REQUIRENAME

377

This tag indicates the dependencies for this package.

378

RPMTAG\_REQUIREVERSION

379

This tag indicates the versions associated with the values found in the RPMTAG\_REQUIRENAME Index.

381

RPMTAG\_CONFLICTFLAGS

382

Bits(s) to specify the conflict range and context.

383

RPMTAG\_CONFLICTNAME

384

This tag indicates the conflicting dependencies for this package.

385

RPMTAG\_CONFLICTVERSION

386

This tag indicates the versions associated with the values found in the RPMTAG\_CONFLICTNAME Index.

388

RPMTAG\_OBSOLETENAME

389

This tag indicates the obsoleted dependencies for this package.

390

RPMTAG\_PROVIDEFLAGS

391

Bits(s) to specify the conflict range and context.

392

RPMTAG\_PROVIDEVERSION

393

This tag indicates the versions associated with the values found in the RPMTAG\_PROVIDENAME Index.

395

RPMTAG\_OBSOLETEFLAGS

396

Bits(s) to specify the conflict range and context.

397

RPMTAG\_OBSOLETEREVISION

398

This tag indicates the versions associated with the values found in the RPMTAG\_OBSOLETENAME Index.

400

#### 22.2.4.4.1 Package Dependency Values

401

The package dependencies are stored in the RPMTAG\_REQUIRENAME and RPMTAG\_REQUIREVERSION index records. The following values may be used.

402

403

**Table 22-13 Index Type values**

Name	Version	Meaning	Status
rpmlib(Versioned Dependencies)	3.0.3-1	Indicates that the package contains RPMTAG_PROVIDENAME, RPMTAG_OBSOLETE_NAME or RPMTAG_PREREQ records that have a version associated with them.	Optional
rpmlib(PayloadFilesHavePrefix)	4.0-1	Indicates the filenames in the Archive have had ":" prepended to them.	Optional
rpmlib(CompressedFileNames)	3.0.4-1	Indicates that the filenames in the Payload are represented in the RPMTAG_DIRINDEXES, RPMTAG_DIRNAME and RPMTAG_BASENAME indexes.	Optional
/bin/sh		Interpreter usually required for installation scripts.	Optional

404

405 Additional dependencies are specified in the Package Dependencies section of this  
 406 specification, and the architecture specific supplements.

407

#### 22.2.4.4.2 Package Dependencies Attributes

408

409 The package dependency attributes are stored in the RPMTAG\_REQUIREFLAGS,  
 410 RPMTAG\_PROVIDEFLAGS and RPMTAG\_OBSOLETEFLAGS index records. The following  
 values may be used.

411

**Table 22-14 Package Dependency Attributes**

Name	Value	Meaning
RPMSENSE_LESS	0x02	
RPMSENSE_GREATER	0x04	
RPMSENSE_EQUAL	0x08	
RPMSENSE_PREREQ	0x40	

Name	Value	Meaning
RPMSENSE_INTERP	0x100	
RPMSENSE_SCRIPT_PRE	0x200	
RPMSENSE_SCRIPT_POST	0x400	
RPMSENSE_SCRIPT_PRUN	0x800	
RPMSENSE_SCRIPT_POSTUN	0x1000	
RPMSENSE_RPMLIB	0x1000000	

412

#### 22.2.4.5 Other Information

413

The following tag values are also found in the Header section.

414

**Table 22-15 Other Tag Values**

Name	Tag Value	Type	Count	Status
RPMTAG_BUILDTIME	1006	INT32	1	Informational
RPMTAG_BUILDHOST	1007	STRING	1	Informational
RPMTAG_FILEVERIFYFLAGS	1045	INT32		Optional
RPMTAG_CHANNELLOGTIME	1080	INT32		Optional
RPMTAG_CHANNELLOGNAME	1081	STRING_ARRAY		Optional
RPMTAG_CHANNELLOGTEXT	1082	STRING_ARRAY		Optional
RPMTAG_OPTFLAGS	1122	STRING	1	Informational
RPMTAG_RHNPLATFORMAT	1131	STRING	1	Deprecated
RPMTAG_PLATFORMAT	1132	STRING	1	Informational

415

RPMTAG\_BUILDTIME

416

This tag specifies the time as seconds since the epoch at which the package was built.

417

RPMTAG\_BUILDHOST

418

This tag specifies the hostname of the system on which the package was built.

419

RPMTAG\_FILEVERIFYFLAGS

420

This tag specifies the bit(s) to control how files are to be verified after install, specifying which checks should be performed.

421

422

423

424

425

```

426 RPMTAG_CHANGELOGTIME
427 This tag specifies the Unix time in seconds since the epoch associated with each
428 entry in the Changelog file.

429 RPMTAG_CHANGELOGNAME
430 This tag specifies the name of who made a change to this package

431 RPMTAG_CHANGELOGTEXT
432 This tag specifies the changes assosciated with a changelog entry.

433 RPMTAG_OPTFLAGS
434 This tag indicates additional flags which may have been passed to the compiler
435 when building this package.

436 RPMTAG_RHNPLATFORM
437 This tag contains an opaque string whose contents are undefined.

438 RPMTAG_PLATFORM
439 This tag contains an opaque string whose contents are undefined.

```

### 22.2.5 Payload Section

```

440 The Payload section contains a compressed cpio archive. The format of this section is
441 defined by RFC 1952: GZIP File Format Specification.
442 When uncompressed, the cpio archive contains a sequence of records for each file.
443 Each record contains a CPIO Header, Filename, Padding, and File Data.

```

**Table 22-16 CPIO File Format**

CPIO Header	Header structure as defined below.
Filename	NUL terminated ASCII string containing the name of the file.
Padding	0-3 bytes as needed to align the file stream to a 4 byte boundary.
File data	The contents of the file.
Padding	0-3 bytes as needed to align the file stream to a 4 byte boundary.

```

445 The CPIO Header uses the following header structure (sometimes referred to as
446 "new ASCII" or "SVR4 cpio"). All numbers are stored as ASCII representations of
447 their hexadecimal value with leading zeros as needed to fill the field. With the
448 exception of c_namesize and the corresponding name string, and c_checksum, all
449 information contained in the CPIO Header is also represented in the Header Section.
450 The values in the CPIO Header shall match the values contained in the Header
451 Section.

```

```

453 struct {
454     char    c_magic[6];
455     char    c_ino[8];
456     char    c_mode[8];
457     char    c_uid[8];

```

```

458          char    c_gid[8];
459          char    c_nlink[8];
460          char    c_mtime[8];
461          char    c_filesize[8];
462          char    c_devmajor[8];
463          char    c_devminor[8];
464          char    c_rdevmajor[8];
465          char    c_rdevminor[8];
466          char    c_namesize[8];
467          char    c_checksum[8];
468      };

469 c_magic
470     Value identifying this cpio format. This value shall be "070701".

471 c_ino
472     This field contains the inode number from the filesystem from which the file
473     was read. This field is ignored when installing a package. This field shall match
474     the corresponding value in the RPMTAG_FILEINODES index in the Header
475     section.

476 c_mode
477     Permission bits of the file. This is an ascii representation of the hexadecimal
478     number representing the bit as defined for the st_mode field of the stat
479     structure defined for the stat function. This field shall match the corresponding
480     value in the RPMTAG_FILEMODES index in the Header section.

481 c_uid
482     Value identifying this owner of this file. This value matches the uid value of the
483     corresponding user in the RPMTAG_FILEUSERNAME as found on the system
484     where this package was built. The username specified in
485     RPMTAG_FILEUSERNAME should take precedence when installing the
486     package.

487 c_gid
488     Value identifying this group of this file. This value matches the gid value of the
489     corresponding user in the RPMTAG_FILEGROUPNAME as found on the
490     system where this package was built. The groupname specified in
491     RPMTAG_FILEGROUPNAME should take precedence when installing the
492     package.

493 c_nlink
494     Value identifying the number of links associated with this file. If the value is
495     greater than 1, then this filename will be linked to 1 or more files in this archive
496     that has a matching value for the c_ino, c_devmajor and c_devminor fields.

497 c_mtime
498     Value identifying the modification time of the file when it was read. This field
499     shall match the corresponding value in the RPMTAG_FILEMTIMES index in the
500     Header section.

501 c_filesize
502     Value identifying the size of the file. This field shall match the corresponding
503     value in the RPMTAG_FILESIZES index in the Header section.

```

504           *c\_devmajor*  
 505           The major number of the device containing the file system from which the file  
 506           was read. With the exception of processing files with *c\_nlink* >1, this field is  
 507           ignored when installing a package. This field shall match the corresponding  
 508           value in the RPMTAG\_FILEDEVICES index in the Header section.

509           *c\_devminor*  
 510           The minor number of the device containing the file system from which the file  
 511           was read. With the exception of processing files with *c\_nlink* >1, this field is  
 512           ignored when installing a package. This field shall match the corresponding  
 513           value in the RPMTAG\_FILEDEVICES index in the Header section.

514           *c\_rdevmajor*  
 515           The major number of the raw device containing the file system from which the file  
 516           was read. This field is ignored when installing a package. This field shall  
 517           match the corresponding value in the RPMTAG\_RDEVS index in the Header  
 518           section.

519           *c\_rdevminor*  
 520           The minor number of the raw device containing the file system from which the file  
 521           was read. This field is ignored when installing a package. This field shall  
 522           match the corresponding value in the RPMTAG\_RDEVS index in the Header  
 523           section.

524           *c\_namesize*  
 525           Value identifying the length of the filename, which is located immediately  
 526           following the CPIO Header structure.

527           *c\_checksum*  
 528           Value containing the CRC checksum of the file data. This field is not used, and  
 529           shall contain the value "00000000". This field is ignored when installing a  
 530           package.

531           A record with the filename "TRAILER!!!" indicates the last record in the archive.

## 22.3 Package Script Restrictions

532           Scripts used as part of the package install and uninstall shall only use commands  
 533           and interfaces that are specified by the LSB. All other commands are not guaranteed  
 534           to be present, or to behave in expected ways.

535           Packages shall not use RPM triggers.

536           Packages shall not depend on the order in which scripts are executed (pre-install,  
 537           pre-uninstall, etc), when doing an upgrade.

## 22.4 Package Tools

538           The LSB does not specify the interface to the tools used to manipulate  
 539           LSB-conformant packages. Each conforming implementation shall provide  
 540           documentation for installing LSB packages.

## 22.5 Package Naming

541 Packages supplied by implementations and applications shall follow the following  
 542 rules for the name field within the package. These rules are not required for the  
 543 filename of the package file itself.

544 **Note:** There are discrepancies among implementations concerning whether the name  
 545 might be `frobnicator-1.7-21-ppc32.rpm` or `frobnicator-1.7-21-powerpc32.rpm`.  
 546 The architecture aside, recommended practice is for the filename of the package file to  
 547 match the name within the package.

548 The following rules apply to the name field alone, not including any release or  
 549 version.

550 **Note:** If the name with the release and version is `frobnicator-1.7-21`, the name part is  
 551 `frobnicator` and falls under the rules for a name with no hyphens.

- 552 • If the name begins with `lsb-` and contains no other hyphens, the name shall be  
 553 assigned by the Linux Assigned Names and Numbers Authority  
 554 (<http://www.lanana.org>) (LANANA), which shall maintain a registry of LSB  
 555 names. The name may be registered by either an implementation or an  
 556 application.
- 557 • If the package name begins with `lsb-` and contains more than one hyphen (for  
 558 example `lsb-distro.example.com-database` or `lsb-gnome-gnumeric`), then  
 559 the portion of the package name between first and second hyphens shall either be  
 560 an LSB provider name assigned by the LANANA, or it may be one of the owners'  
 561 fully-qualified domain names in lower case (e.g., `debian.org`,  
 562 `staroffice.sun.com`). The LSB provider name assigned by LANANA shall only  
 563 consist of the ASCII characters [a-z0-9]. The provider name or domain name may  
 564 be either that of an implementation or an application.
- 565 • Package names containing no hyphens are reserved for use by implementations.  
 566 Applications shall not use such names.
- 567 • Package names which do not start with `lsb-` and which contain a hyphen are  
 568 open to both implementations and applications. Implementations may name  
 569 packages in any part of this namespace. They are encouraged to use names from  
 570 one of the other namespaces available to them, but this is not required due to the  
 571 large amount of current practice to the contrary.

572 **Note:** Widespread existing practice includes such names as `ssh-common`, `ssh-client`,  
 573 `kernel-pcmcia`, and the like. Possible alternative names include `sshcommon`,  
 574 `foolinux-ssh-common` (where `foolinux` is registered to the implementation), or  
 575 `lsb-foolinux-ssh-common`.

576 Applications may name their packages this way, but only if the portion of the  
 577 name before the first hyphen is a provider name or registered domain name as  
 578 described above.

579 **Note:** If an application vendor has domain name such as `visicalc.example.com` and  
 580 has registered `visicalc` as a provider name, they might name packages  
 581 `visicalc-base`, `visicalc.example.com-charting`, and the like.

582 Package names in this namespace are available to both the implementation and an  
 583 application. Implementations and applications will need to consider this potential for  
 584 conflicts when deciding to use these names rather than the alternatives (such as names  
 585 starting with `lsb-`).

## 22.6 Package Dependencies

586 Packages shall have a dependency that indicates which LSB modules are required.  
587 LSB module descriptions are dash seperated tuples containing the name 'lsb', the  
588 module name, and the architecture name. The following dependencies may be used.

589 *lsb-core-arch*

590 This dependency is used to indicate that the application is dependent on  
591 features contained in the LSB-Core specification.

592 *lsb-core-noarch*

593 This dependency is used to indicate that the application is dependent on  
594 features contained in the LSB-Core specification and that the package does not  
595 contain any architecture specific files.

596 These dependencies shall have a version of 3.0.

597 Packages shall not depend on other system-provided dependencies. They shall not  
598 depend on non-system-provided dependencies unless those dependencies are  
599 fulfilled by packages which are part of the same application. A package may only  
600 provide a virtual package name which is registered to that application.

601 Other modules in the LSB may supplement this list. The architecture specific  
602 dependencies are described in the relevant architecture specific LSB.

## 22.7 Package Architecture Considerations

603 Packages which do not contain any architecture specific files should specify an  
604 architecture of noarch. An LSB runtime environment shall accept values noarch, or  
605 the value specified in the architecture specific supplement.

606 Additional specifications or restrictions may be found in the architecture specific  
607 LSB specification.

## Annex A Alphabetical Listing of Interfaces

### A.1 libc

1 The behavior of the interfaces in this library is specified by the following Standards.

Large File Support [LFS]  
This Specification [LSB]  
SUSv2 [SUSv2]  
ISO POSIX (2003) [SUSv3]  
SVID Issue 3 [SVID.3]  
2 SVID Issue 4 [SVID.4]

3 **Table A-1 libc Function Interfaces**

_Exit(GLIBC_2.1.1)[SUSv3]	getpwuid_r(GLIBC_2.1.1)[SUSv3]	sigaddset(GLIBC_2.1.1)[SUSv3]
_IO_feof(GLIBC_2.0)[LSB]	getrlimit(GLIBC_2.0)[SUSv3]	sigaltstack(GLIBC_2.0)[SUSv3]
_IO_getc(GLIBC_2.0)[LSB]	getrlimit64(GLIBC_2.0)[LSB]	sigandset(GLIBC_2.0)[LSB]
_IO_putc(GLIBC_2.0)[LSB]	getrusage(GLIBC_2.0)[SUSv3]	sigdelset(GLIBC_2.0)[SUSv3]
_IO_puts(GLIBC_2.0)[LSB]	getservbyname(GLIBC_2.0)[SUSv3]	sigemptyset(GLIBC_2.0)[SUSv3]
__assert_fail(GLIBC_2.0)[LSB]	getservbyport(GLIBC_2.0)[SUSv3]	sigfillset(GLIBC_2.0)[SUSv3]
__ctype_b_loc[LSB]	getservent()[SUSv3]	sighold()[SUSv3]
__ctype_get_mb_cur_max(GLIBC_2.0)[LSB]	getsid(GLIBC_2.0)[SUSv3]	sigignore(GLIBC_2.0)[SUSv3]
__ctype_tolower_loc[LSB]	getsockname()[SUSv3]	siginterrupt()[SUSv3]
__ctype_toupper_loc[LSB]	getsockopt()[LSB]	sigisemptyset()[LSB]
__cxa_atexit(GLIBC_2.1.3)[LSB]	getsockopt(GLIBC_2.1.3)[SUSv3]	sigismember(GLIBC_2.1.3)[SUSv3]
__errno_location(GLIBC_2.0)[LSB]	gettext(GLIBC_2.0)[LSB]	siglongjmp(GLIBC_2.0)[SUSv3]
__fpending(GLIBC_2.2)[LSB]	gettimeofday(GLIBC_2.2)[SUSv3]	signal(GLIBC_2.2)[SUSv3]
__fxstat(GLIBC_2.0)[LSB]	getuid(GLIBC_2.0)[SUSv3]	sigorset(GLIBC_2.0)[LSB]
__fxstat64(GLIBC_2.2)[LSB]	getutent(GLIBC_2.2)[LSB]	sigpause(GLIBC_2.2)[SUSv3]

<code>__getpagesize(GLIBC_2.0)[LSB]</code>	<code>getutent_r(GLIBC_2.0)[LSB]</code>	<code>sigpending(GLIBC_2.0)[SUSv3]</code>
<code>__getpgid(GLIBC_2.0)[LSB]</code>	<code>getutxent(GLIBC_2.0)[SUSv3]</code>	<code>sigprocmask(GLIBC_2.0)[SUSv3]</code>
<code>__h_errno_location[LSB]</code>	<code>getutxid()[SUSv3]</code>	<code>sigqueue()[SUSv3]</code>
<code>__isinf[LSB]</code>	<code>getutxline()[SUSv3]</code>	<code>sigrelse()[SUSv3]</code>
<code>__isinff[LSB]</code>	<code>getw()[SUSv2]</code>	<code>sigreturn()[LSB]</code>
<code>__isinfl[LSB]</code>	<code>getwc()[SUSv3]</code>	<code>sigset()[SUSv3]</code>
<code>__isnan[LSB]</code>	<code>getwchar()[SUSv3]</code>	<code>sigsuspend()[SUSv3]</code>
<code>__isnanf[LSB]</code>	<code>getwd()[SUSv3]</code>	<code>sigtimedwait()[SUSv3]</code>
<code>__isnanl[LSB]</code>	<code>glob()[SUSv3]</code>	<code>sigwait()[SUSv3]</code>
<code>__libc_current_sigrtmax(GLIBC_2.1)[LSB]</code>	<code>glob64(GLIBC_2.1)[LSB]</code>	<code>sigwaitinfo(GLIBC_2.1)[SUSv3]</code>
<code>__libc_current_sigrtmin(GLIBC_2.1)[LSB]</code>	<code>globfree(GLIBC_2.1)[SUSv3]</code>	<code>sleep(GLIBC_2.1)[SUSv3]</code>
<code>__libc_start_main(GLIBC_2.0)[LSB]</code>	<code>globfree64(GLIBC_2.0)[LSB]</code>	<code>snprintf(GLIBC_2.0)[SUSv3]</code>
<code>__lxstat(GLIBC_2.0)[LSB]</code>	<code>gmtime(GLIBC_2.0)[SUSv3]</code>	<code>sockatmark[SUSv3]</code>
<code>__lxstat64(GLIBC_2.2)[LSB]</code>	<code>gmtime_r(GLIBC_2.2)[SUSv3]</code>	<code>socket(GLIBC_2.2)[SUSv3]</code>
<code>__mempcpy(GLIBC_2.0)[LSB]</code>	<code>grantpt(GLIBC_2.0)[SUSv3]</code>	<code>socketpair(GLIBC_2.0)[SUSv3]</code>
<code>__rawmemchr(GLIBC_2.1)[LSB]</code>	<code>hcreate(GLIBC_2.1)[SUSv3]</code>	<code>sprintf(GLIBC_2.1)[SUSv3]</code>
<code>__register_atfork[LSB]</code>	<code>hdestroy()[SUSv3]</code>	<code>rand()[SUSv3]</code>
<code>__sigsetjmp(GLIBC_2.0)[LSB]</code>	<code>hsearch(GLIBC_2.0)[SUSv3]</code>	<code>rand48(GLIBC_2.0)[SUSv3]</code>
<code>__stpcpy(GLIBC_2.0)[LSB]</code>	<code>htonl(GLIBC_2.0)[SUSv3]</code>	<code>srandom(GLIBC_2.0)[SUSv3]</code>
<code>__strup( GLIBC_2.0)[LSB]</code>	<code>htons(GLIBC_2.0)[SUSv3]</code>	<code>sscanf(GLIBC_2.0)[LSB]</code>
<code>__strtod_internal(GLIBC_2.0)[LSB]</code>	<code>iconv(GLIBC_2.0)[SUSv3]</code>	<code>statvfs(GLIBC_2.0)[SUSv3]</code>
<code>__strtof_internal(GLIBC_2.0)[LSB]</code>	<code>iconv_close(GLIBC_2.0)[SUSv3]</code>	<code>statvfs64[LFS]</code>
<code>__strtok_r(GLIBC_2.0)[LSB]</code>	<code>iconv_open(GLIBC_2.0)[SUSv3]</code>	<code>stime(GLIBC_2.0)[LSB]</code>
<code>__strtol_internal(GLIBC_</code>	<code>if_freetenameindex[SUSv3]</code>	<code>stpcpy(GLIBC_2.0)[LSB]</code>

*Annex A Alphabetical Listing of Interfaces*

2.0)[LSB]		
__strtold_internal(GLIBC_2.0)[LSB]	if_indextoname[SUSv3]	stpncpy(GLIBC_2.0)[LSB]
__strtoll_internal(GLIBC_2.0)[LSB]	if_nameindex[SUSv3]	strcasecmp(GLIBC_2.0)[SUSv3]
__strtoul_internal(GLIBC_2.0)[LSB]	if_nametoindex[SUSv3]	strcasestr(GLIBC_2.0)[LSB]
__strtoull_internal(GLIBC_2.0)[LSB]	imaxabs(GLIBC_2.0)[SUSv3]	strcat(GLIBC_2.0)[SUSv3]
__sysconf(GLIBC_2.2)[LSB]	imaxdiv(GLIBC_2.2)[SUSv3]	strchr(GLIBC_2.2)[SUSv3]
__sysv_signal(GLIBC_2.0)[LSB]	index(GLIBC_2.0)[SUSv3]	strcmp(GLIBC_2.0)[SUSv3]
__wcstod_internal(GLIBC_2.0)[LSB]	inet_addr(GLIBC_2.0)[SUSv3]	strcoll(GLIBC_2.0)[SUSv3]
__wcstof_internal(GLIBC_2.0)[LSB]	inet_ntoa(GLIBC_2.0)[SUSv3]	strcpy(GLIBC_2.0)[SUSv3]
__wcstol_internal(GLIBC_2.0)[LSB]	inet_ntop[SUSv3]	strcspn(GLIBC_2.0)[SUSv3]
__wcstold_internal(GLIBC_2.0)[LSB]	inet_pton[SUSv3]	strupr(GLIBC_2.0)[SUSv3]
__wcstoul_internal(GLIBC_2.0)[LSB]	initgroups(GLIBC_2.0)[LSB]	strerror(GLIBC_2.0)[SUSv3]
__xmknod(GLIBC_2.0)[LSB]	initstate(GLIBC_2.0)[SUSv3]	strerror_r(GLIBC_2.0)[LSB]
__xstat(GLIBC_2.0)[LSB]	insque(GLIBC_2.0)[SUSv3]	strfmon(GLIBC_2.0)[SUSv3]
__xstat64(GLIBC_2.2)[LSB]	ioctl(GLIBC_2.2)[LSB]	strftime(GLIBC_2.2)[SUSv3]
_exit(GLIBC_2.0)[SUSv3]	isalnum(GLIBC_2.0)[SUSv3]	strlen(GLIBC_2.0)[SUSv3]
_longjmp(GLIBC_2.0)[SUSv3]	isalpha(GLIBC_2.0)[SUSv3]	strncasecmp(GLIBC_2.0)[SUSv3]
_setjmp(GLIBC_2.0)[SUSv3]	isascii(GLIBC_2.0)[SUSv3]	strncat(GLIBC_2.0)[SUSv3]
_tolower(GLIBC_2.0)[SUSv3]	isatty(GLIBC_2.0)[SUSv3]	strncmp(GLIBC_2.0)[SUSv3]
_toupper(GLIBC_2.0)[SUSv3]	isblank(GLIBC_2.0)[SUSv3]	strncpy(GLIBC_2.0)[SUSv3]
a64l(GLIBC_2.0)[SUSv3]	iscntrl(GLIBC_2.0)[SUSv3]	strndup(GLIBC_2.0)[LSB]

abort(GLIBC_2.0)[SUSv3]	isdigit(GLIBC_2.0)[SUSv3]	strnlen(GLIBC_2.0)[LSB]
abs(GLIBC_2.0)[SUSv3]	isgraph(GLIBC_2.0)[SUSv3]	strupbrk(GLIBC_2.0)[SUSv3]
accept(GLIBC_2.0)[SUSv3]	islower(GLIBC_2.0)[SUSv3]	strptime(GLIBC_2.0)[LSB]
access(GLIBC_2.0)[SUSv3]	isprint(GLIBC_2.0)[SUSv3]	strrchr(GLIBC_2.0)[SUSv3]
acct(GLIBC_2.0)[LSB]	ispunct(GLIBC_2.0)[SUSv3]	strsep(GLIBC_2.0)[LSB]
adjtime(GLIBC_2.0)[LSB]	isspace(GLIBC_2.0)[SUSv3]	strsignal(GLIBC_2.0)[LSB]
alarm(GLIBC_2.0)[SUSv3]	isupper(GLIBC_2.0)[SUSv3]	strspn(GLIBC_2.0)[SUSv3]
asctime(GLIBC_2.0)[SUSv3]	iswalnum(GLIBC_2.0)[SUSv3]	strstr(GLIBC_2.0)[SUSv3]
asctime_r(GLIBC_2.0)[SUSv3]	iswalpha(GLIBC_2.0)[SUSv3]	strtod(GLIBC_2.0)[SUSv3]
asprintf(GLIBC_2.0)[LSB]	iswblank(GLIBC_2.0)[SUSv3]	strtof(GLIBC_2.0)[SUSv3]
atof(GLIBC_2.0)[SUSv3]	iswcntrl(GLIBC_2.0)[SUSv3]	strtoimax(GLIBC_2.0)[SUSv3]
atoi(GLIBC_2.0)[SUSv3]	iswctype(GLIBC_2.0)[SUSv3]	strtok(GLIBC_2.0)[SUSv3]
atol(GLIBC_2.0)[SUSv3]	iswdigit(GLIBC_2.0)[SUSv3]	strtok_r(GLIBC_2.0)[SUSv3]
atoll[SUSv3]	iswgraph()[SUSv3]	strtol()[SUSv3]
authnone_create(GLIBC_2.0)[SVID.4]	iswlower(GLIBC_2.0)[SUSv3]	strtold(GLIBC_2.0)[SUSv3]
basename(GLIBC_2.0)[SUSv3]	iswprint(GLIBC_2.0)[SUSv3]	strtoll(GLIBC_2.0)[SUSv3]
bcmp(GLIBC_2.0)[SUSv3]	iswpunct(GLIBC_2.0)[SUSv3]	strtoq(GLIBC_2.0)[LSB]
bcopy(GLIBC_2.0)[SUSv3]	iswspace(GLIBC_2.0)[SUSv3]	strtoul(GLIBC_2.0)[SUSv3]
bind(GLIBC_2.0)[SUSv3]	iswupper(GLIBC_2.0)[SUSv3]	strtoull(GLIBC_2.0)[SUSv3]
bind_textdomain_codeset[LSB]	iswxdigit()[SUSv3]	strtoumax()[SUSv3]
bindresvport(GLIBC_2.0)[LSB]	isxdigit(GLIBC_2.0)[SUSv3]	strtouq(GLIBC_2.0)[LSB]

*Annex A Alphabetical Listing of Interfaces*

bindtextdomain(GLIBC_2.0)[LSB]	jrand48(GLIBC_2.0)[SUSv3]	strxfrm(GLIBC_2.0)[SUSv3]
brk(GLIBC_2.0)[SUSv2]	key_decryptsession(GLIBC_2.0)[SVID.3]	svc_getreqset(GLIBC_2.0)[SVID.3]
bsd_signal(GLIBC_2.0)[SUSv3]	kill(GLIBC_2.0)[LSB]	svc_register(GLIBC_2.0)[LSB]
bsearch(GLIBC_2.0)[SUSv3]	killpg(GLIBC_2.0)[SUSv3]	svc_run(GLIBC_2.0)[LSB]
btowc(GLIBC_2.0)[SUSv3]	l64a(GLIBC_2.0)[SUSv3]	svc_sendreply(GLIBC_2.0)[LSB]
bzero(GLIBC_2.0)[SUSv3]	labs(GLIBC_2.0)[SUSv3]	svcerr_auth(GLIBC_2.0)[SVID.3]
calloc(GLIBC_2.0)[SUSv3]	lchown(GLIBC_2.0)[SUSv3]	svcerr_decode(GLIBC_2.0)[SVID.3]
catclose(GLIBC_2.0)[SUSv3]	lcong48(GLIBC_2.0)[SUSv3]	svcerr_noproc(GLIBC_2.0)[SVID.3]
catgets(GLIBC_2.0)[SUSv3]	ldiv(GLIBC_2.0)[SUSv3]	svcerr_noprog(GLIBC_2.0)[SVID.3]
catopen(GLIBC_2.0)[SUSv3]	lfind(GLIBC_2.0)[SUSv3]	svcerr_progvers(GLIBC_2.0)[SVID.3]
cfgetispeed(GLIBC_2.0)[SUSv3]	link(GLIBC_2.0)[LSB]	svcerr_systemerr(GLIBC_2.0)[SVID.3]
cfgetospeed(GLIBC_2.0)[SUSv3]	listen(GLIBC_2.0)[SUSv3]	svcerr_weakauth(GLIBC_2.0)[SVID.3]
cfmakeraw(GLIBC_2.0)[LSB]	llabs(GLIBC_2.0)[SUSv3]	svctcp_create(GLIBC_2.0)[LSB]
cfsetispeed(GLIBC_2.0)[SUSv3]	lldiv(GLIBC_2.0)[SUSv3]	svcudp_create(GLIBC_2.0)[LSB]
cfsetospeed(GLIBC_2.0)[SUSv3]	localeconv(GLIBC_2.0)[SUSv3]	swab(GLIBC_2.0)[SUSv3]
cfsetspeed(GLIBC_2.0)[LSB]	localtime(GLIBC_2.0)[SUSv3]	swapcontext(GLIBC_2.0)[SUSv3]
chdir(GLIBC_2.0)[SUSv3]	localtime_r(GLIBC_2.0)[SUSv3]	swprintf(GLIBC_2.0)[SUSv3]
chmod(GLIBC_2.0)[SUSv3]	lockf(GLIBC_2.0)[SUSv3]	swscanf(GLIBC_2.0)[LSB]
chown(GLIBC_2.1)[SUSv3]	lockf64(GLIBC_2.1)[LFS]	symlink(GLIBC_2.1)[SUSv3]
chroot(GLIBC_2.0)[SUSv2]	longjmp(GLIBC_2.0)[SUSv3]	sync(GLIBC_2.0)[SUSv3]
clearerr(GLIBC_2.0)[SUS]	lrand48(GLIBC_2.0)[SUS]	sysconf(GLIBC_2.0)[SUS]

v3]	v3]	v3]
clnt_create(GLIBC_2.0)[S VID.4]	lsearch(GLIBC_2.0)[SUSv 3]	syslog(GLIBC_2.0)[SUSv 3]
clnt_pcreateerror(GLIBC _2.0)[SVID.4]	lseek(GLIBC_2.0)[SUSv3]	system(GLIBC_2.0)[LSB]
clnt_perrno(GLIBC_2.0)[ SVID.4]	makecontext(GLIBC_2.0) [SUSv3]	tcdrain(GLIBC_2.0)[SUS v3]
clnt_perror(GLIBC_2.0)[S VID.4]	malloc(GLIBC_2.0)[SUSv 3]	tcflow(GLIBC_2.0)[SUSv 3]
clnt_spcreateerror(GLIB C_2.0)[SVID.4]	mblen(GLIBC_2.0)[SUSv 3]	tcflush(GLIBC_2.0)[SUSv 3]
clnt_sperrno(GLIBC_2.0) [SVID.4]	mbrlen(GLIBC_2.0)[SUS v3]	tcgetattr(GLIBC_2.0)[SUS v3]
clnt_sperror(GLIBC_2.0)[ SVID.4]	mbrtowc(GLIBC_2.0)[SU Sv3]	tcgetpgrp(GLIBC_2.0)[S USv3]
clock(GLIBC_2.0)[SUSv3]	mbsinit(GLIBC_2.0)[SUS v3]	tcgetsid(GLIBC_2.0)[SUS v3]
close(GLIBC_2.0)[SUSv3]	mbsnr towcs(GLIBC_2.0)[ LSB]	tcsendbreak(GLIBC_2.0)[ SUSv3]
closedir(GLIBC_2.0)[SUS v3]	mbsrtowcs(GLIBC_2.0)[S USv3]	tcsetattr(GLIBC_2.0)[SUS v3]
closelog(GLIBC_2.0)[SUS v3]	mbstowcs(GLIBC_2.0)[S USv3]	tcsetpgrp(GLIBC_2.0)[S USv3]
confstr(GLIBC_2.0)[SUSv 3]	mbtowc(GLIBC_2.0)[SUS v3]	tdelete[SUSv3]
connect(GLIBC_2.0)[SUS v3]	memccpy(GLIBC_2.0)[S USv3]	telldir(GLIBC_2.0)[SUSv 3]
creat(GLIBC_2.0)[SUSv3]	memchr(GLIBC_2.0)[SUS v3]	tempnam(GLIBC_2.0)[S USv3]
creat64(GLIBC_2.1)[LFS]	memcmp(GLIBC_2.1)[SU Sv3]	textdomain(GLIBC_2.1)[ LSB]
ctermid(GLIBC_2.0)[SUS v3]	memcpy(GLIBC_2.0)[SU Sv3]	tfind(GLIBC_2.0)[SUSv3]
ctime(GLIBC_2.0)[SUSv3 ]	memmem(GLIBC_2.0)[L SB]	time(GLIBC_2.0)[SUSv3]
ctime_r(GLIBC_2.0)[SUS v3]	memmove(GLIBC_2.0)[S USv3]	times(GLIBC_2.0)[SUSv3 ]
cuserid(GLIBC_2.0)[SUS v2]	memrchr(GLIBC_2.0)[LS B]	tmpfile(GLIBC_2.0)[SUS v3]
daemon(GLIBC_2.0)[LSB ]	memset(GLIBC_2.0)[SUS v3]	tmpfile64(GLIBC_2.0)[LF S]

*Annex A Alphabetical Listing of Interfaces*

dcgettext(GLIBC_2.0)[LSB]	mkdir(GLIBC_2.0)[SUSv3]	tmpnam(GLIBC_2.0)[SUSv3]
dcngettext[LSB]	mkfifo()[SUSv3]	toascii()[SUSv3]
dgettext[LSB]	mkstemp()[SUSv3]	tolower()[SUSv3]
difftime(GLIBC_2.0)[SUSv3]	mkstemp64(GLIBC_2.0)[LFS]	toupper(GLIBC_2.0)[SUSv3]
dirname(GLIBC_2.0)[SUSv3]	mktemp(GLIBC_2.0)[SUSv3]	towctrans(GLIBC_2.0)[SUSv3]
div(GLIBC_2.0)[SUSv3]	mktimes(GLIBC_2.0)[SUSv3]	towlower(GLIBC_2.0)[SUSv3]
dnggettext[LSB]	mlock()[SUSv3]	towupper()[SUSv3]
drand48(GLIBC_2.0)[SUSv3]	mlockall(GLIBC_2.0)[SUSv3]	truncate(GLIBC_2.0)[SUSv3]
dup(GLIBC_2.0)[SUSv3]	mmap(GLIBC_2.0)[SUSv3]	truncate64(GLIBC_2.0)[LFS]
dup2(GLIBC_2.0)[SUSv3]	mmap64(GLIBC_2.0)[LF S]	tsearch(GLIBC_2.0)[SUSv3]
duplocale[LSB]	mprotect()[SUSv3]	ttyname()[SUSv3]
ecvt(GLIBC_2.0)[SUSv3]	mrand48(GLIBC_2.0)[SUSv3]	ttyname_r(GLIBC_2.0)[SUSv3]
endgrent(GLIBC_2.0)[SUSv3]	msgctl(GLIBC_2.0)[SUSv3]	twalk(GLIBC_2.0)[SUSv3]
endprotoent(GLIBC_2.0)[SUSv3]	msgget(GLIBC_2.0)[SUSv3]	tzset(GLIBC_2.0)[SUSv3]
endpwent(GLIBC_2.0)[SUSv3]	msgrcv(GLIBC_2.0)[SUSv3]	ualarm(GLIBC_2.0)[SUSv3]
endservent(GLIBC_2.0)[SUSv3]	msgsnd(GLIBC_2.0)[SUSv3]	ulimit(GLIBC_2.0)[SUSv3]
endutent(GLIBC_2.0)[SUSv2]	msync(GLIBC_2.0)[SUSv3]	umask(GLIBC_2.0)[SUSv3]
endutxent(GLIBC_2.1)[SUSv3]	munlock(GLIBC_2.1)[SUSv3]	uname(GLIBC_2.1)[SUSv3]
erand48(GLIBC_2.0)[SUSv3]	munlockall(GLIBC_2.0)[SUSv3]	ungetc(GLIBC_2.0)[SUSv3]
err(GLIBC_2.0)[LSB]	munmap(GLIBC_2.0)[SUSv3]	ungetwc(GLIBC_2.0)[SUSv3]
error(GLIBC_2.0)[LSB]	nanosleep(GLIBC_2.0)[SUSv3]	unlink(GLIBC_2.0)[LSB]
errx(GLIBC_2.0)[LSB]	newlocale[LSB]	unlockpt(GLIBC_2.0)[SUSv3]

execl(GLIBC_2.0)[SUSv3]	nftw(GLIBC_2.0)[SUSv3]	unsetenv[SUSv3]
execle(GLIBC_2.0)[SUSv3]	nftw64(GLIBC_2.0)[LFS]	uselocale[LSB]
execlp(GLIBC_2.0)[SUSv3]	gettext[LSB]	usleep(GLIBC_2.0)[SUSv3]
execv(GLIBC_2.0)[SUSv3]	nice(GLIBC_2.0)[SUSv3]	utime(GLIBC_2.0)[SUSv3]
execve(GLIBC_2.0)[SUSv3]	nl_langinfo(GLIBC_2.0)[SUSv3]	utimes(GLIBC_2.0)[SUSv3]
execvp(GLIBC_2.0)[SUSv3]	rand48(GLIBC_2.0)[SUSv3]	utmpname[LSB]
exit(GLIBC_2.0)[SUSv3]	ntohl(GLIBC_2.0)[SUSv3]	vasprintf(GLIBC_2.0)[LSB]
fchdir(GLIBC_2.0)[SUSv3]	ntohs(GLIBC_2.0)[SUSv3]	vdprintf(GLIBC_2.0)[LSB]
fchmod(GLIBC_2.0)[SUSv3]	open(GLIBC_2.0)[SUSv3]	verrx(GLIBC_2.0)[LSB]
fchown(GLIBC_2.0)[SUSv3]	opendir(GLIBC_2.0)[SUSv3]	vfork(GLIBC_2.0)[SUSv3]
fclose(GLIBC_2.1)[SUSv3]	openlog(GLIBC_2.1)[SUSv3]	vfprintf(GLIBC_2.1)[SUSv3]
fcntl(GLIBC_2.0)[LSB]	pathconf(GLIBC_2.0)[SUSv3]	vfscanf[LSB]
fcvt(GLIBC_2.0)[SUSv3]	pause(GLIBC_2.0)[SUSv3]	vfwprintf(GLIBC_2.0)[SUSv3]
fdatasync(GLIBC_2.0)[SUSv3]	pclose(GLIBC_2.0)[SUSv3]	vfwscanf(GLIBC_2.0)[LSB]
fdopen(GLIBC_2.1)[SUSv3]	perror(GLIBC_2.1)[SUSv3]	vprintf(GLIBC_2.1)[SUSv3]
feof(GLIBC_2.0)[SUSv3]	pipe(GLIBC_2.0)[SUSv3]	vscanf[LSB]
ferror(GLIBC_2.0)[SUSv3]	pmap_getport(GLIBC_2.0)[LSB]	vsnprintf(GLIBC_2.0)[SUSv3]
fflush(GLIBC_2.0)[SUSv3]	pmap_set(GLIBC_2.0)[LSB]	vsprintf(GLIBC_2.0)[SUSv3]
fflush_unlocked(GLIBC_2.0)[LSB]	pmap_unset(GLIBC_2.0)[LSB]	vsscanf[LSB]
ffs(GLIBC_2.0)[SUSv3]	poll(GLIBC_2.0)[SUSv3]	vswprintf(GLIBC_2.0)[SUSv3]
fgetc(GLIBC_2.0)[SUSv3]	popen(GLIBC_2.0)[SUSv3]	vswscanf(GLIBC_2.0)[LSB]

*Annex A Alphabetical Listing of Interfaces*

fgetpos(GLIBC_2.0)[SUSv3]	posix_memalign(GLIBC_2.0)[SUSv3]	vsyslog[LSB]
fgetpos64(GLIBC_2.1)[LFS]	posix_openpt[SUSv3]	vwprintf(GLIBC_2.1)[SUSv3]
fgets(GLIBC_2.0)[SUSv3]	printf(GLIBC_2.0)[SUSv3]	vwscanf(GLIBC_2.0)[LSB]
fgetwc(GLIBC_2.2)[SUSv3]	psignal(GLIBC_2.2)[LSB]	wait(GLIBC_2.2)[SUSv3]
fgetwc_unlocked(GLIBC_2.2)[LSB]	ptsname(GLIBC_2.2)[SUSv3]	wait4(GLIBC_2.2)[LSB]
fgetws(GLIBC_2.2)[SUSv3]	putc(GLIBC_2.2)[SUSv3]	waitpid(GLIBC_2.2)[LSB]
fileno(GLIBC_2.0)[SUSv3]	putc_unlocked(GLIBC_2.0)[SUSv3]	warn(GLIBC_2.0)[LSB]
flock(GLIBC_2.0)[LSB]	putchar(GLIBC_2.0)[SUSv3]	warnx(GLIBC_2.0)[LSB]
flockfile(GLIBC_2.0)[SUSv3]	putchar_unlocked(GLIBC_2.0)[SUSv3]	wcpncpy(GLIBC_2.0)[LSB]
fmtmsg(GLIBC_2.1)[SUSv3]	putenv(GLIBC_2.1)[SUSv3]	wcpncpy(GLIBC_2.1)[LSB]
fnmatch(GLIBC_2.2.3)[SUSv3]	puts(GLIBC_2.2.3)[SUSv3]	wcrtomb(GLIBC_2.2.3)[SUSv3]
fopen(GLIBC_2.1)[SUSv3]	pututxline(GLIBC_2.1)[SUSv3]	wcscasecmp(GLIBC_2.1)[LSB]
fopen64(GLIBC_2.1)[LFS]	putw(GLIBC_2.1)[SUSv2]	wcscat(GLIBC_2.1)[SUSv3]
fork(GLIBC_2.0)[SUSv3]	putwc(GLIBC_2.0)[SUSv3]	wcschr(GLIBC_2.0)[SUSv3]
fpathconf(GLIBC_2.0)[SUSv3]	putwchar(GLIBC_2.0)[SUSv3]	wcscmp(GLIBC_2.0)[SUSv3]
fprintf(GLIBC_2.0)[SUSv3]	qsort(GLIBC_2.0)[SUSv3]	wcscoll(GLIBC_2.0)[SUSv3]
fputc(GLIBC_2.0)[SUSv3]	raise(GLIBC_2.0)[SUSv3]	wcscpy(GLIBC_2.0)[SUSv3]
fputs(GLIBC_2.0)[SUSv3]	rand(GLIBC_2.0)[SUSv3]	wcscspn(GLIBC_2.0)[SUSv3]
fputwc(GLIBC_2.2)[SUSv3]	rand_r(GLIBC_2.2)[SUSv3]	wcsdup(GLIBC_2.2)[LSB]
fputws(GLIBC_2.2)[SUSv3]	random(GLIBC_2.2)[SUSv3]	wcsftime(GLIBC_2.2)[SUSv3]
fread(GLIBC_2.0)[SUSv3]	read(GLIBC_2.0)[SUSv3]	wcslen(GLIBC_2.0)[SUSv3]

]		3]
free(GLIBC_2.0)[SUSv3]	readdir(GLIBC_2.0)[SUSv3]	wcsncasecmp(GLIBC_2.0)[LSB]
freeaddrinfo[SUSv3]	readdir64()[LFS]	wcsncat()[SUSv3]
freelocale[LSB]	readdir_r[SUSv3]	wcsncmp()[SUSv3]
freopen(GLIBC_2.0)[SUSv3]	readlink(GLIBC_2.0)[SUSv3]	wcsncpy(GLIBC_2.0)[SUSv3]
freopen64(GLIBC_2.1)[LFS]	readv(GLIBC_2.1)[SUSv3]	wcsnlen(GLIBC_2.1)[LSB]
fscanf(GLIBC_2.0)[LSB]	realloc(GLIBC_2.0)[SUSv3]	wcsnrombs(GLIBC_2.0)[LSB]
fseek(GLIBC_2.0)[SUSv3]	realpath(GLIBC_2.0)[SUSv3]	wcspbrk(GLIBC_2.0)[SUSv3]
fseeko(GLIBC_2.1)[SUSv3]	recv(GLIBC_2.1)[SUSv3]	wcsrchr(GLIBC_2.1)[SUSv3]
fseeko64(GLIBC_2.1)[LFS]	recvfrom(GLIBC_2.1)[SUSv3]	wcsrtombs(GLIBC_2.1)[SUSv3]
fsetpos(GLIBC_2.0)[SUSv3]	recvmsg(GLIBC_2.0)[SUSv3]	wcsspn(GLIBC_2.0)[SUSv3]
fsetpos64(GLIBC_2.1)[LFS]	regcomp(GLIBC_2.1)[SUSv3]	wcsstr(GLIBC_2.1)[SUSv3]
fstatvfs(GLIBC_2.1)[SUSv3]	regerror(GLIBC_2.1)[SUSv3]	wcstod(GLIBC_2.1)[SUSv3]
fstatvfs64(GLIBC_2.1)[LFS]	regexec(GLIBC_2.1)[LSB]	wcstof(GLIBC_2.1)[SUSv3]
fsync(GLIBC_2.0)[SUSv3]	regfree(GLIBC_2.0)[SUSv3]	wcstoi(max)(GLIBC_2.0)[SUSv3]
ftell(GLIBC_2.0)[SUSv3]	remove(GLIBC_2.0)[SUSv3]	wcstok(GLIBC_2.0)[SUSv3]
ftello(GLIBC_2.1)[SUSv3]	remque(GLIBC_2.1)[SUSv3]	wcstol(GLIBC_2.1)[SUSv3]
ftello64(GLIBC_2.1)[LFS]	rename(GLIBC_2.1)[SUSv3]	wcstold(GLIBC_2.1)[SUSv3]
ftime(GLIBC_2.0)[SUSv3]	rewind(GLIBC_2.0)[SUSv3]	wcstoll(GLIBC_2.0)[SUSv3]
ftok(GLIBC_2.0)[SUSv3]	rewinddir(GLIBC_2.0)[SUSv3]	wcstombs(GLIBC_2.0)[SUSv3]
ftruncate(GLIBC_2.0)[SUSv3]	rindex(GLIBC_2.0)[SUSv3]	wcstoq(GLIBC_2.0)[LSB]
ftruncate64(GLIBC_2.1)[	rmdir(GLIBC_2.1)[SUSv3]	wcstoul(GLIBC_2.1)[SUS

*Annex A Alphabetical Listing of Interfaces*

LFS]	]	v3]
ftrylockfile(GLIBC_2.0)[SUSv3]	sbrk(GLIBC_2.0)[SUSv2]	wcstoull(GLIBC_2.0)[SUSv3]
ftw(GLIBC_2.0)[SUSv3]	scanf(GLIBC_2.0)[LSB]	wcstoumax(GLIBC_2.0)[SUSv3]
ftw64(GLIBC_2.1)[LFS]	sched_get_priority_max(GLIBC_2.1)[SUSv3]	wcstouq(GLIBC_2.1)[LSB]
funlockfile(GLIBC_2.0)[SUSv3]	sched_get_priority_min(GLIBC_2.0)[SUSv3]	wcswcs(GLIBC_2.0)[SUSv3]
fwide(GLIBC_2.2)[SUSv3]	sched_getparam(GLIBC_2.2)[SUSv3]	wcswidth(GLIBC_2.2)[SUSv3]
fwprintf(GLIBC_2.2)[SUSv3]	sched_getscheduler(GLIBC_2.2)[SUSv3]	wcsxfrm(GLIBC_2.2)[SUSv3]
fwrite(GLIBC_2.0)[SUSv3]	sched_rr_get_interval(GLIBC_2.0)[SUSv3]	wctob(GLIBC_2.0)[SUSv3]
fwscanf(GLIBC_2.2)[LSB]	sched_setparam(GLIBC_2.2)[SUSv3]	wctomb(GLIBC_2.2)[SUSv3]
gai_strerror[SUSv3]	sched_setscheduler()[SUSv3]	wctrans()[SUSv3]
gcvt(GLIBC_2.0)[SUSv3]	sched_yield(GLIBC_2.0)[SUSv3]	wctype(GLIBC_2.0)[SUSv3]
getaddrinfo[SUSv3]	seed48()[SUSv3]	wcwidth()[SUSv3]
getc(GLIBC_2.0)[SUSv3]	seekdir(GLIBC_2.0)[SUSv3]	wmemchr(GLIBC_2.0)[SUSv3]
getc_unlocked(GLIBC_2.0)[SUSv3]	select(GLIBC_2.0)[SUSv3]	wmemcmp(GLIBC_2.0)[SUSv3]
getchar(GLIBC_2.0)[SUSv3]	semctl(GLIBC_2.0)[SUSv3]	wmemcpy(GLIBC_2.0)[SUSv3]
getchar_unlocked(GLIBC_2.0)[SUSv3]	semget(GLIBC_2.0)[SUSv3]	wmemmove(GLIBC_2.0)[SUSv3]
getcontext(GLIBC_2.1)[SUSv3]	semop(GLIBC_2.1)[SUSv3]	wmemset(GLIBC_2.1)[SUSv3]
getcwd(GLIBC_2.0)[SUSv3]	send(GLIBC_2.0)[SUSv3]	wordexp(GLIBC_2.0)[SUSv3]
getdate(GLIBC_2.1)[SUSv3]	sendmsg(GLIBC_2.1)[SUSv3]	wordfree(GLIBC_2.1)[SUSv3]
getegid(GLIBC_2.0)[SUSv3]	sendto(GLIBC_2.0)[SUSv3]	wprintf(GLIBC_2.0)[SUSv3]
getenv(GLIBC_2.0)[SUSv3]	setbuf(GLIBC_2.0)[SUSv3]	write(GLIBC_2.0)[SUSv3]
geteuid(GLIBC_2.0)[SUS]	setbuffer(GLIBC_2.0)[LS]	writev(GLIBC_2.0)[SUSv]

v3]	B]	3]
getgid(GLIBC_2.0)[SUSv3]	setcontext(GLIBC_2.0)[SUSv3]	wscanf(GLIBC_2.0)[LSB]
getgrent(GLIBC_2.0)[SUSv3]	setegid(GLIBC_2.0)[SUSv3]	xdr_accepted_reply(GLIBC_2.0)[SVID.3]
getgrgid(GLIBC_2.0)[SUSv3]	setenv[SUSv3]	xdr_array(GLIBC_2.0)[SVID.3]
getgrgid_r(GLIBC_2.0)[SUSv3]	seteuid(GLIBC_2.0)[SUSv3]	xdr_bool(GLIBC_2.0)[SVID.3]
getgrnam(GLIBC_2.0)[SUSv3]	setgid(GLIBC_2.0)[SUSv3]	xdr_bytes(GLIBC_2.0)[SVID.3]
getgrnam_r(GLIBC_2.0)[SUSv3]	setgrent(GLIBC_2.0)[SUSv3]	xdr_callhdr(GLIBC_2.0)[SVID.3]
getgrouplist[LSB]	setgroups()[LSB]	xdr_callmsg()[SVID.3]
getgroups(GLIBC_2.0)[SUSv3]	sethostname(GLIBC_2.0)[LSB]	xdr_char(GLIBC_2.0)[SVID.3]
gethostbyaddr(GLIBC_2.0)[SUSv3]	setitimer(GLIBC_2.0)[SUSv3]	xdr_double(GLIBC_2.0)[SVID.3]
gethostbyname(GLIBC_2.0)[SUSv3]	setlocale(GLIBC_2.0)[SUSv3]	xdr_enum(GLIBC_2.0)[SVID.3]
gethostid(GLIBC_2.0)[SUSv3]	setlogmask(GLIBC_2.0)[SUSv3]	xdr_float(GLIBC_2.0)[SVID.3]
gethostname(GLIBC_2.0)[SUSv3]	setpgid(GLIBC_2.0)[SUSv3]	xdr_free(GLIBC_2.0)[SVID.3]
getitimer(GLIBC_2.0)[SUSv3]	setpgrp(GLIBC_2.0)[SUSv3]	xdr_int(GLIBC_2.0)[SVID.3]
getloadavg(GLIBC_2.2)[LSB]	setpriority(GLIBC_2.2)[SUSv3]	xdr_long(GLIBC_2.2)[SVID.3]
getlogin(GLIBC_2.0)[SUSv3]	setprotoent(GLIBC_2.0)[SUSv3]	xdr_opaque(GLIBC_2.0)[SVID.3]
getlogin_r[SUSv3]	setpwent()[SUSv3]	xdr_opaque_auth()[SVID.3]
getnameinfo[SUSv3]	setregid()[SUSv3]	xdr_pointer()[SVID.3]
getopt(GLIBC_2.0)[LSB]	setreuid(GLIBC_2.0)[SUSv3]	xdr_reference(GLIBC_2.0)[SVID.3]
getopt_long(GLIBC_2.0)[LSB]	setrlimit(GLIBC_2.0)[SUSv3]	xdr_rejected_reply(GLIBC_2.0)[SVID.3]
getopt_long_only(GLIBC_2.0)[LSB]	setrlimit64[LFS]	xdr_repliesmsg(GLIBC_2.0)[SVID.3]
getpagesize(GLIBC_2.0)[	setservent(GLIBC_2.0)[S	xdr_short(GLIBC_2.0)[S

SUSv2]	USv3]	VID.3]
getpeername(GLIBC_2.0)[SUSv3]	setsid(GLIBC_2.0)[SUSv3]	xdr_string(GLIBC_2.0)[S VID.3]
getpgid(GLIBC_2.0)[SUSv3]	setsockopt(GLIBC_2.0)[LSB]	xdr_u_char(GLIBC_2.0)[ SVID.3]
getpgrp(GLIBC_2.0)[SUSv3]	setstate(GLIBC_2.0)[SUSv3]	xdr_u_int(GLIBC_2.0)[LSB]
getpid(GLIBC_2.0)[SUSv3]	setuid(GLIBC_2.0)[SUSv3]	xdr_u_long(GLIBC_2.0)[ SVID.3]
getppid(GLIBC_2.0)[SUSv3]	setutent(GLIBC_2.0)[LSB]	xdr_u_short(GLIBC_2.0)[ SVID.3]
getpriority(GLIBC_2.0)[SUSv3]	setutxent(GLIBC_2.0)[SUSv3]	xdr_union(GLIBC_2.0)[S VID.3]
getprotobyname(GLIBC_2.0)[SUSv3]	setvbuf(GLIBC_2.0)[SUSv3]	xdr_vector(GLIBC_2.0)[S VID.3]
getprotobynumber(GLIBC_2.0)[SUSv3]	shmat(GLIBC_2.0)[SUSv3]	xdr_void(GLIBC_2.0)[S VID.3]
getprotoent(GLIBC_2.0)[SUSv3]	shmctl(GLIBC_2.0)[SUSv3]	xdr_wrapstring(GLIBC_2. .0)[SVID.3]
getpwent(GLIBC_2.0)[SUSv3]	shmdt(GLIBC_2.0)[SUSv3]	xdrmem_create(GLIBC_2. .0)[SVID.3]
getpwnam(GLIBC_2.0)[SUSv3]	shmget(GLIBC_2.0)[SUSv3]	xdrrec_create(GLIBC_2.0) )[SVID.3]
getpwnam_r(GLIBC_2.0)[SUSv3]	shutdown(GLIBC_2.0)[SUSv3]	xdrrec_eof(GLIBC_2.0)[S VID.3]
getpwuid(GLIBC_2.0)[SUSv3]	sigaction(GLIBC_2.0)[SUSv3]	

4

5

**Table A-2 libc Data Interfaces**

<u>__daylight</u> <a href="#">ID_STD_46</a> <u>LSB</u>	<u>__timezone</u> <a href="#">ID_STD_46</a> <u>LSB</u>	<u>__sys_errlist</u> <a href="#">ID_STD_46</a> <u>LSB</u>
<u>__environ</u> <a href="#">ID_STD_46</a> <u>LSB</u>	<u>__tzname</u> <a href="#">ID_STD_46</a> <u>LSB</u>	

6

## A.2 libcrypt

7

8

The behavior of the interfaces in this library is specified by the following Standards.  
ISO POSIX (2003) [SUSv3]

9

**Table A-3 libcrypt Function Interfaces**

crypt(GLIBC_2.0)[SUSv3]	encrypt(GLIBC_2.0)[SUSv3]	setkey(GLIBC_2.0)[SUSv3]
-------------------------	---------------------------	--------------------------

10

### A.3 libdl

The behavior of the interfaces in this library is specified by the following Standards.  
 This Specification [LSB]  
 ISO POSIX (2003) [SUSv3]

**Table A-4 libdl Function Interfaces**

dladdr(GLIBC_2.0)[LSB]	dlerror(GLIBC_2.0)[SUSv3]	dlsym(GLIBC_2.0)[LSB]
dlclose(GLIBC_2.0)[SUSv3]	dlopen(GLIBC_2.0)[LSB]	

### A.4 libm

The behavior of the interfaces in this library is specified by the following Standards.  
 ISO C (1999) [ISOC99]  
 This Specification [LSB]  
 SUSv2 [SUSv2]  
 ISO POSIX (2003) [SUSv3]

**Table A-5 libm Function Interfaces**

__finite[ISOC99]	csinhf()[SUSv3]	log10()[SUSv3]
__finitef[ISOC99]	csinhl()[SUSv3]	log10f[SUSv3]
__finitel[ISOC99]	csinl()[SUSv3]	log10l[SUSv3]
__fpclassify[LSB]	csqrtf()[SUSv3]	log1p()[SUSv3]
__fpclassifyf[LSB]	csqrftf()[SUSv3]	log1pf[SUSv3]
__signbit[ISOC99]	csqrfl()[SUSv3]	log1pl[SUSv3]
__signbitf[ISOC99]	ctanf()[SUSv3]	log2f[SUSv3]
acos(GLIBC_2.0)[SUSv3]	ctanf(GLIBC_2.0)[SUSv3]	log2f[SUSv3]
acosf(GLIBC_2.0)[SUSv3]	ctanh(GLIBC_2.0)[SUSv3]	log2l[SUSv3]
acosh(GLIBC_2.0)[SUSv3]	ctanhf(GLIBC_2.0)[SUSv3]	logb(GLIBC_2.0)[SUSv3]
acoshf(GLIBC_2.0)[SUSv3]	ctanhlf(GLIBC_2.0)[SUSv3]	logbf[SUSv3]
acoshf(GLIBC_2.0)[SUSv3]	ctanlf(GLIBC_2.0)[SUSv3]	logbl[SUSv3]
acosl(GLIBC_2.0)[SUSv3]	dremf(GLIBC_2.0)[ISOC99]	logf[SUSv3]
asin(GLIBC_2.0)[SUSv3]	dremlf(GLIBC_2.0)[ISOC99]	logl[SUSv3]
asinf(GLIBC_2.0)[SUSv3]	erf(GLIBC_2.0)[SUSv3]	lrint(GLIBC_2.0)[SUSv3]

*Annex A Alphabetical Listing of Interfaces*

asinh(GLIBC_2.0)[SUSv3]	erfc(GLIBC_2.0)[SUSv3]	lrintf(GLIBC_2.0)[SUSv3]
asinhf(GLIBC_2.0)[SUSv3]	erfcf(GLIBC_2.0)[SUSv3]	lrintl(GLIBC_2.0)[SUSv3]
asinhl(GLIBC_2.0)[SUSv3]	erfc1(GLIBC_2.0)[SUSv3]	lround(GLIBC_2.0)[SUSv3]
asinl(GLIBC_2.0)[SUSv3]	erff(GLIBC_2.0)[SUSv3]	lroundf(GLIBC_2.0)[SUSv3]
atan(GLIBC_2.0)[SUSv3]	erfl(GLIBC_2.0)[SUSv3]	lroundl(GLIBC_2.0)[SUSv3]
atan2(GLIBC_2.0)[SUSv3]	exp(GLIBC_2.0)[SUSv3]	matherr(GLIBC_2.0)[ISO C99]
atan2f(GLIBC_2.0)[SUSv3]	exp2[SUSv3]	modf(GLIBC_2.0)[SUSv3]
atan2l(GLIBC_2.0)[SUSv3]	exp2f[SUSv3]	modff(GLIBC_2.0)[SUSv3]
atanf(GLIBC_2.0)[SUSv3]	expf[SUSv3]	modfl(GLIBC_2.0)[SUSv3]
atanh(GLIBC_2.0)[SUSv3]	expl[SUSv3]	nan(GLIBC_2.0)[SUSv3]
atanhf(GLIBC_2.0)[SUSv3]	expm1(GLIBC_2.0)[SUSv3]	nanf(GLIBC_2.0)[SUSv3]
atanhl(GLIBC_2.0)[SUSv3]	expm1f[SUSv3]	nanl(GLIBC_2.0)[SUSv3]
atanl(GLIBC_2.0)[SUSv3]	expm1l[SUSv3]	nearbyint(GLIBC_2.0)[SUSv3]
cabs(GLIBC_2.1)[SUSv3]	fabs(GLIBC_2.1)[SUSv3]	nearbyintf(GLIBC_2.1)[SUSv3]
cabsf(GLIBC_2.1)[SUSv3]	fabsf(GLIBC_2.1)[SUSv3]	nearbyintl(GLIBC_2.1)[SUSv3]
cabsl(GLIBC_2.1)[SUSv3]	fabsl(GLIBC_2.1)[SUSv3]	nextafter(GLIBC_2.1)[SUSv3]
cacos(GLIBC_2.1)[SUSv3]	fdim(GLIBC_2.1)[SUSv3]	nextafterf(GLIBC_2.1)[SUSv3]
cacosf(GLIBC_2.1)[SUSv3]	fdimf(GLIBC_2.1)[SUSv3]	nextafterl(GLIBC_2.1)[SUSv3]
cacosh(GLIBC_2.1)[SUSv3]	fdiml(GLIBC_2.1)[SUSv3]	nexttoward(GLIBC_2.1)[SUSv3]
cacoshf(GLIBC_2.1)[SUSv3]	feclearexcept(GLIBC_2.1)[SUSv3]	nexttowardf(GLIBC_2.1)[SUSv3]
cacoshl(GLIBC_2.1)[SUS]	fegetenv(GLIBC_2.1)[SU]	nexttowardl(GLIBC_2.1)[

v3]	Sv3]	SUSv3]
cacosl(GLIBC_2.1)[SUSv3]	fegetexceptflag(GLIBC_2.1)[SUSv3]	pow(GLIBC_2.1)[SUSv3]
carg(GLIBC_2.1)[SUSv3]	fegetround(GLIBC_2.1)[SUSv3]	pow10(GLIBC_2.1)[ISO C99]
cargf(GLIBC_2.1)[SUSv3]	feholdexcept(GLIBC_2.1)[SUSv3]	pow10f(GLIBC_2.1)[ISO C99]
cargl(GLIBC_2.1)[SUSv3]	feraiseexcept(GLIBC_2.1)[SUSv3]	pow10l(GLIBC_2.1)[ISO C99]
casin(GLIBC_2.1)[SUSv3]	fesetenv(GLIBC_2.1)[SUSv3]	powf(GLIBC_2.1)[SUSv3]
casinf(GLIBC_2.1)[SUSv3]	fesetexceptflag(GLIBC_2.1)[SUSv3]	powl(GLIBC_2.1)[SUSv3]
casinh(GLIBC_2.1)[SUSv3]	fesetround(GLIBC_2.1)[SUSv3]	remainder(GLIBC_2.1)[SUSv3]
casinhf(GLIBC_2.1)[SUSv3]	fetestexcept(GLIBC_2.1)[SUSv3]	remainderf(GLIBC_2.1)[SUSv3]
casinhl(GLIBC_2.1)[SUSv3]	feupdateenv(GLIBC_2.1)[SUSv3]	remainderl(GLIBC_2.1)[SUSv3]
casinl(GLIBC_2.1)[SUSv3]	finite(GLIBC_2.1)[SUSv2]	remquo(GLIBC_2.1)[SUSv3]
catan(GLIBC_2.1)[SUSv3]	finitef(GLIBC_2.1)[ISO C99]	remquof(GLIBC_2.1)[SUSv3]
catanf(GLIBC_2.1)[SUSv3]	finitel(GLIBC_2.1)[ISO C99]	remquol(GLIBC_2.1)[SUSv3]
catanh(GLIBC_2.1)[SUSv3]	floor(GLIBC_2.1)[SUSv3]	rint(GLIBC_2.1)[SUSv3]
catanhf(GLIBC_2.1)[SUSv3]	floorf(GLIBC_2.1)[SUSv3]	rintf(GLIBC_2.1)[SUSv3]
catanhhl(GLIBC_2.1)[SUSv3]	floorl(GLIBC_2.1)[SUSv3]	rintl(GLIBC_2.1)[SUSv3]
catanl(GLIBC_2.1)[SUSv3]	fma(GLIBC_2.1)[SUSv3]	round(GLIBC_2.1)[SUSv3]
cbrt(GLIBC_2.0)[SUSv3]	fmaf(GLIBC_2.0)[SUSv3]	roundf(GLIBC_2.0)[SUSv3]
cbrtf(GLIBC_2.0)[SUSv3]	fmal(GLIBC_2.0)[SUSv3]	roundl(GLIBC_2.0)[SUSv3]
cbrtl(GLIBC_2.0)[SUSv3]	fmax(GLIBC_2.0)[SUSv3]	scalb(GLIBC_2.0)[SUSv3]
ccos(GLIBC_2.1)[SUSv3]	fmaxf(GLIBC_2.1)[SUSv3]	scalbf(GLIBC_2.1)[ISO C99]

*Annex A Alphabetical Listing of Interfaces*

ccosf(GLIBC_2.1)[SUSv3]	fmaxl(GLIBC_2.1)[SUSv3]	scalbl(GLIBC_2.1)[ISOC99]
ccosh(GLIBC_2.1)[SUSv3]	fmin(GLIBC_2.1)[SUSv3]	scalbln(GLIBC_2.1)[SUSv3]
ccoshf(GLIBC_2.1)[SUSv3]	fminf(GLIBC_2.1)[SUSv3]	scalblnf(GLIBC_2.1)[SUSv3]
ccoshl(GLIBC_2.1)[SUSv3]	fminl(GLIBC_2.1)[SUSv3]	scalblnl(GLIBC_2.1)[SUSv3]
ccosl(GLIBC_2.1)[SUSv3]	fmod(GLIBC_2.1)[SUSv3]	scalbn(GLIBC_2.1)[SUSv3]
ceil(GLIBC_2.0)[SUSv3]	fmodf(GLIBC_2.0)[SUSv3]	scalbnf(GLIBC_2.0)[SUSv3]
ceilf(GLIBC_2.0)[SUSv3]	fmodl(GLIBC_2.0)[SUSv3]	scalbnl(GLIBC_2.0)[SUSv3]
ceill(GLIBC_2.0)[SUSv3]	frexp(GLIBC_2.0)[SUSv3]	significand(GLIBC_2.0)[ISOC99]
cexp(GLIBC_2.1)[SUSv3]	frexpf(GLIBC_2.1)[SUSv3]	significandf(GLIBC_2.1)[ISOC99]
cexpf(GLIBC_2.1)[SUSv3]	frexpl(GLIBC_2.1)[SUSv3]	significandl(GLIBC_2.1)[ISOC99]
cexpl(GLIBC_2.1)[SUSv3]	gamma(GLIBC_2.1)[SUSv2]	sin(GLIBC_2.1)[SUSv3]
cimag(GLIBC_2.1)[SUSv3]	gammaf(GLIBC_2.1)[ISO C99]	sincos(GLIBC_2.1)[ISOC99]
cimaf(GLIBC_2.1)[SUSv3]	gammal(GLIBC_2.1)[ISO C99]	sincosf(GLIBC_2.1)[ISOC99]
cimagl(GLIBC_2.1)[SUSv3]	hypot(GLIBC_2.1)[SUSv3]	sincosl(GLIBC_2.1)[ISOC99]
clog(GLIBC_2.1)[SUSv3]	hypotf(GLIBC_2.1)[SUSv3]	sinf(GLIBC_2.1)[SUSv3]
clog10(GLIBC_2.1)[ISOC99]	hypotl(GLIBC_2.1)[SUSv3]	sinh(GLIBC_2.1)[SUSv3]
clog10f(GLIBC_2.1)[ISO C99]	ilogb(GLIBC_2.1)[SUSv3]	sinhf(GLIBC_2.1)[SUSv3]
clog10l(GLIBC_2.1)[ISOC99]	ilogbf(GLIBC_2.1)[SUSv3]	sinhl(GLIBC_2.1)[SUSv3]
clogf(GLIBC_2.1)[SUSv3]	ilogbl(GLIBC_2.1)[SUSv3]	sinl(GLIBC_2.1)[SUSv3]
clogl(GLIBC_2.1)[SUSv3]	j0(GLIBC_2.1)[SUSv3]	sqrt(GLIBC_2.1)[SUSv3]
conj(GLIBC_2.1)[SUSv3]	j0f(GLIBC_2.1)[ISOC99]	sqrtf(GLIBC_2.1)[SUSv3]

conjf(GLIBC_2.1)[SUSv3]	j0l(GLIBC_2.1)[ISOC99]	sqrtl(GLIBC_2.1)[SUSv3]
conjl(GLIBC_2.1)[SUSv3]	j1(GLIBC_2.1)[SUSv3]	tan(GLIBC_2.1)[SUSv3]
copysign(GLIBC_2.0)[SUSv3]	j1f(GLIBC_2.0)[ISOC99]	tanf(GLIBC_2.0)[SUSv3]
copysignf(GLIBC_2.0)[SUSv3]	j1l(GLIBC_2.0)[ISOC99]	tanh(GLIBC_2.0)[SUSv3]
copysignl(GLIBC_2.0)[SUSv3]	jn(GLIBC_2.0)[SUSv3]	tanhf(GLIBC_2.0)[SUSv3]
cos(GLIBC_2.0)[SUSv3]	jnf(GLIBC_2.0)[ISOC99]	tanhl(GLIBC_2.0)[SUSv3]
cosf(GLIBC_2.0)[SUSv3]	jnl(GLIBC_2.0)[ISOC99]	tanl(GLIBC_2.0)[SUSv3]
cosh(GLIBC_2.0)[SUSv3]	ldexp(GLIBC_2.0)[SUSv3]	tgamma(GLIBC_2.0)[SUSv3]
coshf(GLIBC_2.0)[SUSv3]	ldexpf(GLIBC_2.0)[SUSv3]	tgammaf(GLIBC_2.0)[SUSv3]
coshl(GLIBC_2.0)[SUSv3]	ldexpl(GLIBC_2.0)[SUSv3]	tgammal(GLIBC_2.0)[SUSv3]
cosl(GLIBC_2.0)[SUSv3]	lgamma(GLIBC_2.0)[SUSv3]	trunc(GLIBC_2.0)[SUSv3]
cpow(GLIBC_2.1)[SUSv3]	lgamma_r(GLIBC_2.1)[ISOC99]	truncf(GLIBC_2.1)[SUSv3]
cpowf(GLIBC_2.1)[SUSv3]	lgammaf(GLIBC_2.1)[SUSv3]	truncl(GLIBC_2.1)[SUSv3]
cpowl(GLIBC_2.1)[SUSv3]	lgammaf_r(GLIBC_2.1)[ISOC99]	y0(GLIBC_2.1)[SUSv3]
cproj(GLIBC_2.1)[SUSv3]	lgammal(GLIBC_2.1)[SUSv3]	y0f(GLIBC_2.1)[ISOC99]
cprojf(GLIBC_2.1)[SUSv3]	lgammal_r(GLIBC_2.1)[ISOC99]	y0l(GLIBC_2.1)[ISOC99]
cprojl(GLIBC_2.1)[SUSv3]	llrint(GLIBC_2.1)[SUSv3]	y1(GLIBC_2.1)[SUSv3]
creal(GLIBC_2.1)[SUSv3]	llrintf(GLIBC_2.1)[SUSv3]	y1f(GLIBC_2.1)[ISOC99]
crealf(GLIBC_2.1)[SUSv3]	llrintl(GLIBC_2.1)[SUSv3]	y1l(GLIBC_2.1)[ISOC99]
creall(GLIBC_2.1)[SUSv3]	llround(GLIBC_2.1)[SUSv3]	yn(GLIBC_2.1)[SUSv3]
csin(GLIBC_2.1)[SUSv3]	llroundf(GLIBC_2.1)[SUSv3]	ynf(GLIBC_2.1)[ISOC99]
csinf(GLIBC_2.1)[SUSv3]	llroundl(GLIBC_2.1)[SUSv3]	ynl(GLIBC_2.1)[ISOC99]

18

csinh(GLIBC_2.1)[SUSv3] ]	log(GLIBC_2.1)[SUSv3]	
------------------------------	-----------------------	--

19

**Table A-6 libm Data Interfaces**

20

sign-gamID_STD_46_SUS <a href="#">V3</a>		
---	--	--

## A.5 libncurses

21

The behavior of the interfaces in this library is specified by the following Standards.

22

X/Open Curses [SUS-CURSES]

23

**Table A-7 libncurses Function Interfaces**

addch[SUS-CURSES]	mvdelch[SUS-CURSES]	slk_refresh[SUS-CURSES] ]
addchnstr[SUS-CURSES]	mvderwin[SUS-CURSES]	slk_restore[SUS-CURSES] ]
addchstr[SUS-CURSES]	mvgetch[SUS-CURSES]	slk_set[SUS-CURSES]
addnstr[SUS-CURSES]	mvgetnstr[SUS-CURSES]	slk_touch[SUS-CURSES]
addstr[SUS-CURSES]	mvgetstr[SUS-CURSES]	standend[SUS-CURSES]
attr_get[SUS-CURSES]	mvhline[SUS-CURSES]	standout[SUS-CURSES]
attr_off[SUS-CURSES]	mvinch[SUS-CURSES]	start_color[SUS-CURSES] ]
attr_on[SUS-CURSES]	mvinchnstr[SUS-CURSES]	subpad[SUS-CURSES]
attr_set[SUS-CURSES]	mvinchstr[SUS-CURSES]	subwin[SUS-CURSES]
attroff[SUS-CURSES]	mvinnstr[SUS-CURSES]	syncok[SUS-CURSES]
attron[SUS-CURSES]	mvinsch[SUS-CURSES]	termattrs[SUS-CURSES]
attrset[SUS-CURSES]	mvinsnstr[SUS-CURSES]	termname[SUS-CURSES]
baudrate[SUS-CURSES]	mvinsstr[SUS-CURSES]	tgetent[SUS-CURSES]
beep[SUS-CURSES]	mvinstr[SUS-CURSES]	tgetflag[SUS-CURSES]
bkgd[SUS-CURSES]	mvprintw[SUS-CURSES]	tgetnum[SUS-CURSES]
bkgdset[SUS-CURSES]	mvscanw[SUS-CURSES]	tgetstr[SUS-CURSES]
border[SUS-CURSES]	mvvline[SUS-CURSES]	tgoto[SUS-CURSES]
box[SUS-CURSES]	mvwaddch[SUS-CURSES]	tigetflag[SUS-CURSES]
can_change_color[SUS-CURSES]	mvwaddchnstr[SUS-CURSES]	tigetnum[SUS-CURSES]

cbreak[SUS-CURSES]	mvwaddchstr[SUS-CURSES]	tigetstr[SUS-CURSES]
chgat[SUS-CURSES]	mvwaddnstr[SUS-CURSES]	timeout[SUS-CURSES]
clear[SUS-CURSES]	mvwaddstr[SUS-CURSES]	touchline[SUS-CURSES]
clearok[SUS-CURSES]	mvwchgat[SUS-CURSES]	touchwin[SUS-CURSES]
clrtoobot[SUS-CURSES]	mvwdelch[SUS-CURSES]	tparm[SUS-CURSES]
clrtoeol[SUS-CURSES]	mvwgetch[SUS-CURSES]	tputs[SUS-CURSES]
color_content[SUS-CURSES]	mvwgetnstr[SUS-CURSES]	typeahead[SUS-CURSES]
color_set[SUS-CURSES]	mvwgetstr[SUS-CURSES]	unctrl[SUS-CURSES]
copywin[SUS-CURSES]	mvwhline[SUS-CURSES]	ungetch[SUS-CURSES]
curs_set[SUS-CURSES]	mvwin[SUS-CURSES]	untouchwin[SUS-CURSES]
def_prog_mode[SUS-CURSES]	mvwinch[SUS-CURSES]	use_env[SUS-CURSES]
def_shell_mode[SUS-CURSES]	mvwinchnstr[SUS-CURSES]	vidattr[SUS-CURSES]
del_curterm[SUS-CURSES]	mvwinchstr[SUS-CURSES]	vidputs[SUS-CURSES]
delay_output[SUS-CURSES]	mvwinnstr[SUS-CURSES]	vline[SUS-CURSES]
delch[SUS-CURSES]	mvwinsch[SUS-CURSES]	vwprintw[SUS-CURSES]
deleteln[SUS-CURSES]	mvwinsnstr[SUS-CURSES]	vw_scanw[SUS-CURSES]
delscreen[SUS-CURSES]	mvwinsstr[SUS-CURSES]	vwprintw[SUS-CURSES]
delwin[SUS-CURSES]	mvwinstr[SUS-CURSES]	vwscanw[SUS-CURSES]
derwin[SUS-CURSES]	mvwprintw[SUS-CURSES]	waddch[SUS-CURSES]
doupdate[SUS-CURSES]	mvwscanw[SUS-CURSES]	waddchnstr[SUS-CURSES]
dupwin[SUS-CURSES]	mvwvline[SUS-CURSES]	waddchstr[SUS-CURSES]

*Annex A Alphabetical Listing of Interfaces*

echo[SUS-CURSES]	napms[SUS-CURSES]	waddnstr[SUS-CURSES]
echochar[SUS-CURSES]	newpad[SUS-CURSES]	waddstr[SUS-CURSES]
endwin[SUS-CURSES]	newterm[SUS-CURSES]	wattr_get[SUS-CURSES]
erase[SUS-CURSES]	newwin[SUS-CURSES]	wattr_off[SUS-CURSES]
erasechar[SUS-CURSES]	nl[SUS-CURSES]	wattr_on[SUS-CURSES]
filter[SUS-CURSES]	nocbreak[SUS-CURSES]	wattr_set[SUS-CURSES]
flash[SUS-CURSES]	nodelay[SUS-CURSES]	wattroff[SUS-CURSES]
flushinp[SUS-CURSES]	noecho[SUS-CURSES]	wattron[SUS-CURSES]
getbkgd[SUS-CURSES]	nonl[SUS-CURSES]	wattrset[SUS-CURSES]
getch[SUS-CURSES]	noqiflush[SUS-CURSES]	wbkgd[SUS-CURSES]
getnstr[SUS-CURSES]	noraw[SUS-CURSES]	wbkgdset[SUS-CURSES]
getstr[SUS-CURSES]	notimeout[SUS-CURSES]	wborder[SUS-CURSES]
getwin[SUS-CURSES]	overlay[SUS-CURSES]	wchgat[SUS-CURSES]
halfdelay[SUS-CURSES]	overwrite[SUS-CURSES]	wclear[SUS-CURSES]
has_colors[SUS-CURSES]	pair_content[SUS-CURSES]	wclrtoobot[SUS-CURSES]
has_ic[SUS-CURSES]	pechochar[SUS-CURSES]	wclrtoeol[SUS-CURSES]
has_il[SUS-CURSES]	pnoutrefresh[SUS-CURSES]	wcolor_set[SUS-CURSES]
hline[SUS-CURSES]	prefresh[SUS-CURSES]	wcursyncup[SUS-CURSES]
idcok[SUS-CURSES]	printw[SUS-CURSES]	wdelch[SUS-CURSES]
idlok[SUS-CURSES]	putp[SUS-CURSES]	wdeleteln[SUS-CURSES]
immedok[SUS-CURSES]	putwin[SUS-CURSES]	wechochar[SUS-CURSES]
inch[SUS-CURSES]	qiflush[SUS-CURSES]	werase[SUS-CURSES]
inchnstr[SUS-CURSES]	raw[SUS-CURSES]	wgetch[SUS-CURSES]
inchstr[SUS-CURSES]	redrawwin[SUS-CURSES]	wgetnstr[SUS-CURSES]
init_color[SUS-CURSES]	refresh[SUS-CURSES]	wgetstr[SUS-CURSES]
init_pair[SUS-CURSES]	reset_prog_mode[SUS-CURSES]	whline[SUS-CURSES]
initscr[SUS-CURSES]	reset_shell_mode[SUS-CURSES]	winch[SUS-CURSES]
innstr[SUS-CURSES]	resety[SUS-CURSES]	winchnstr[SUS-CURSES]
insch[SUS-CURSES]	restartterm[SUS-CURSES]	winchstr[SUS-CURSES]

	]	
insdelln[SUS-CURSES]	rioffline[SUS-CURSES]	winnstr[SUS-CURSES]
insertln[SUS-CURSES]	savetty[SUS-CURSES]	winsch[SUS-CURSES]
insnstr[SUS-CURSES]	scanw[SUS-CURSES]	winsdelln[SUS-CURSES]
insstr[SUS-CURSES]	scr_dump[SUS-CURSES]	winsertln[SUS-CURSES]
instr[SUS-CURSES]	scr_init[SUS-CURSES]	winsnstr[SUS-CURSES]
intrflush[SUS-CURSES]	scr_restore[SUS-CURSES] ]	winsstr[SUS-CURSES]
is_linetouched[SUS-CURSES]	scr_set[SUS-CURSES]	winstr[SUS-CURSES]
is_wintouched[SUS-CURSES]	scrl[SUS-CURSES]	wmove[SUS-CURSES]
isendwin[SUS-CURSES]	scroll[SUS-CURSES]	wnoutrefresh[SUS-CURSES]
keyname[SUS-CURSES]	scrolllok[SUS-CURSES]	wprintw[SUS-CURSES]
keypad[SUS-CURSES]	set_curerm[SUS-CURSES]	wredrawln[SUS-CURSES] ]
killchar[SUS-CURSES]	set_term[SUS-CURSES]	wrefresh[SUS-CURSES]
leaveok[SUS-CURSES]	setscreg[SUS-CURSES]	wscanw[SUS-CURSES]
longname[SUS-CURSES]	setupterm[SUS-CURSES]	wscrl[SUS-CURSES]
meta[SUS-CURSES]	slk_attr_set[SUS-CURSES]	wsetscreg[SUS-CURSES] ]
move[SUS-CURSES]	slk_atroff[SUS-CURSES]	wstandend[SUS-CURSES] ]
mvaddch[SUS-CURSES]	slk_attron[SUS-CURSES]	wstandout[SUS-CURSES] ]
mvaddchnstr[SUS-CURSES]	slk_attrset[SUS-CURSES]	wsyncdown[SUS-CURSES]
mvaddchstr[SUS-CURSES]	slk_clear[SUS-CURSES]	wsyncup[SUS-CURSES]
mvaddnstr[SUS-CURSES] ]	slk_color[SUS-CURSES]	wtimeout[SUS-CURSES]
mvaddstr[SUS-CURSES]	slk_init[SUS-CURSES]	wtouchln[SUS-CURSES]
mvchgat[SUS-CURSES]	slk_label[SUS-CURSES]	wvline[SUS-CURSES]
mvcur[SUS-CURSES]	slk_noutrefresh[SUS-CURSES]	

**Table A-8 libncurses Data Interfaces**

COLORS	ID_STD_46_S	LINES	ID_STD_46_SU	curscr	ID_STD_46_SU
--------	-------------	-------	--------------	--------	--------------

26

<u>US_46_CURSES</u>	<u>S_46_CURSES</u>	<u>S_46_CURSES</u>
COLOR_PAIRS <u>ID_STD_46_SUS_46_CURSES_S</u>	acs_map <u>ID_STD_46_SUS_46_CURSES</u>	stdscr <u>ID_STD_46_SUS_46_CURSES</u>
COLS <u>ID_STD_46_SUS_46_CURSES</u>	cur_term <u>ID_STD_46_SUS_46_CURSES</u>	

## A.6 libpam

27

The behavior of the interfaces in this library is specified by the following Standards.

28

This Specification [LSB]

29

**Table A-9 libpam Function Interfaces**

30

pam_acct_mgmt[LSB]	pam_fail_delay[LSB]	pam_setcred[LSB]
pam_authenticate[LSB]	pam_get_item[LSB]	pam_start[LSB]
pam_chauthtok[LSB]	pam_getenvlist[LSB]	pam_strerror[LSB]
pam_close_session[LSB]	pam_open_session[LSB]	
pam_end[LSB]	pam_set_item[LSB]	

## A.7 libpthread

31

The behavior of the interfaces in this library is specified by the following Standards.

32

Large File Support [LFS]

This Specification [LSB]

ISO POSIX (2003) [SUSv3]

33

**Table A-10 pthread Function Interfaces**

_pthread_cleanup_pop[LSB]	pthread_cond_wait()[SUSv3]	pthread_rwlock_timedunlock[USv3]
_pthread_cleanup_push[LSB]	pthread_condattr_destroy()[USv3]	pthread_rwlock_tryrdlock[USv3]
lseek64(GLIBC_2.1)[LFS]	pthread_condattr_getpshared[SUSv3]	pthread_rwlock_trywrlock(GLIBC_2.1)[USv3]
open64(GLIBC_2.1)[LFS]	pthread_condattr_init(GLIBC_2.1)[USv3]	pthread_rwlock_unlock(GLIBC_2.1)[USv3]
pread(GLIBC_2.1)[USv3]	pthread_condattr_setpshared[SUSv3]	pthread_rwlock_wrlock(GLIBC_2.1)[USv3]
pread64(GLIBC_2.1)[LFS]	pthread_create(GLIBC_2.1)[USv3]	pthread_rwlockattr_destroy(GLIBC_2.1)[USv3]
pthread_attr_destroy(GLIBC_2.0)[USv3]	pthread_detach(GLIBC_2.0)[USv3]	pthread_rwlockattr_getpshared(GLIBC_2.0)[USv3]
pthread_attr_getdetachst	pthread_equal(GLIBC_2.	pthread_rwlockattr_init(

ate(GLIBC_2.0)[SUSv3]	0)[SUSv3]	GLIBC_2.0)[SUSv3]
pthread_attr_getguardsize(GLIBC_2.1)[SUSv3]	pthread_exit(GLIBC_2.1)[SUSv3]	pthread_rwlockattr_setpshared(GLIBC_2.1)[SUSv3]
pthread_attr_getinheritsched(GLIBC_2.0)[SUSv3]	pthread_getconcurrency[SUSv3]	pthread_self(GLIBC_2.0)[SUSv3]
pthread_attr_getschedparam(GLIBC_2.0)[SUSv3]	pthread_getschedparam(GLIBC_2.0)[SUSv3]	pthread_setcancelstate(GLIBC_2.0)[SUSv3]
pthread_attr_getschedpolicy(GLIBC_2.0)[SUSv3]	pthread_getspecific(GLIBC_2.0)[SUSv3]	pthread_setcanceltype(GLIBC_2.0)[SUSv3]
pthread_attr_getscope(GLIBC_2.0)[SUSv3]	pthread_join(GLIBC_2.0)[SUSv3]	pthread_setconcurrency[SUSv3]
pthread_attr_getstack[SUSv3]	pthread_key_create()[SUSv3]	pthread_setschedparam()[SUSv3]
pthread_attr_getstackaddr(GLIBC_2.1)[SUSv3]	pthread_key_delete(GLIBC_2.1)[SUSv3]	pthread_setschedprio[SUSv3]
pthread_attr_getstacksize(GLIBC_2.1)[SUSv3]	pthread_kill(GLIBC_2.1)[SUSv3]	pthread_setspecific(GLIBC_2.1)[SUSv3]
pthread_attr_init(GLIBC_2.1)[SUSv3]	pthread_mutex_destroy(GLIBC_2.1)[SUSv3]	pthread_sigmask(GLIBC_2.1)[SUSv3]
pthread_attr_setdetachstate(GLIBC_2.0)[SUSv3]	pthread_mutex_init(GLIBC_2.0)[SUSv3]	pthread_testcancel(GLIBC_2.0)[SUSv3]
pthread_attr_setguardsize(GLIBC_2.1)[SUSv3]	pthread_mutex_lock(GLIBC_2.1)[SUSv3]	pwrite(GLIBC_2.1)[SUSv3]
pthread_attr_setinheritsched(GLIBC_2.0)[SUSv3]	pthread_mutex_trylock(GLIBC_2.0)[SUSv3]	pwrite64(GLIBC_2.0)[LFS]
pthread_attr_setschedparam(GLIBC_2.0)[SUSv3]	pthread_mutex_unlock(GLIBC_2.0)[SUSv3]	sem_close(GLIBC_2.0)[SUSv3]
pthread_attr_setschedpolicy(GLIBC_2.0)[SUSv3]	pthread_mutexattr_destr oy(GLIBC_2.0)[SUSv3]	sem_destroy(GLIBC_2.0)[SUSv3]
pthread_attr_setscope(GLIBC_2.0)[SUSv3]	pthread_mutexattr_getshared(GLIBC_2.0)[SUSv3]	sem_getvalue(GLIBC_2.0)[SUSv3]
pthread_attr_setstack[SUSv3]	pthread_mutexattr_getty pe()[SUSv3]	sem_init()[SUSv3]
pthread_attr_setstackaddr(GLIBC_2.1)[SUSv3]	pthread_mutexattr_init(GLIBC_2.1)[SUSv3]	sem_open(GLIBC_2.1)[SUSv3]
pthread_attr_setstacksize(GLIBC_2.1)[SUSv3]	pthread_mutexattr_setp shared(GLIBC_2.1)[SUSv3]	sem_post(GLIBC_2.1)[SUSv3]
pthread_cancel(GLIBC_2.0)[SUSv3]	pthread_mutexattr_settymeout(GLIBC_2.0)[SUSv3]	sem_timedwait(GLIBC_2.0)[SUSv3]

34

pthread_cond_broadcast(GLIBC_2.0)[SUSv3]	pthread_once(GLIBC_2.0)[SUSv3]	sem_trywait(GLIBC_2.0)[SUSv3]
pthread_cond_destroy(GLIBC_2.0)[SUSv3]	pthread_rwlock_destroy(GLIBC_2.0)[SUSv3]	sem_unlink(GLIBC_2.0)[SUSv3]
pthread_cond_init(GLIBC_2.0)[SUSv3]	pthread_rwlock_init(GLIBC_2.0)[SUSv3]	sem_wait(GLIBC_2.0)[SUSv3]
pthread_cond_signal(GLIBC_2.0)[SUSv3]	pthread_rwlock_rdlock(GLIBC_2.0)[SUSv3]	
pthread_cond_timedwait(GLIBC_2.0)[SUSv3]	pthread_rwlock_timedrdlock[SUSv3]	

## A.8 librt

35

The behavior of the interfaces in this library is specified by the following Standards.

36

ISO POSIX (2003) [SUSv3]

37

**Table A-11 librt Function Interfaces**

38

clock_getcpu_clockid(GLIBC_2.2)[SUSv3]	clock_settime(GLIBC_2.2)[SUSv3]	timer_delete(GLIBC_2.2)[SUSv3]
clock_getres(GLIBC_2.2)[SUSv3]	shm_open(GLIBC_2.2)[SUSv3]	timer_getoverrun(GLIBC_2.2)[SUSv3]
clock_gettime(GLIBC_2.2)[SUSv3]	shm_unlink(GLIBC_2.2)[SUSv3]	timer_gettime(GLIBC_2.2)[SUSv3]
clock_nanosleep(GLIBC_2.2)[SUSv3]	timer_create(GLIBC_2.2)[SUSv3]	timer_settime(GLIBC_2.2)[SUSv3]

## A.9 libutil

39

The behavior of the interfaces in this library is specified by the following Standards.

40

This Specification [LSB]

41

**Table A-12 libutil Function Interfaces**

42

forkpty(GLIBC_2.0)[LSB]	login_tty(GLIBC_2.0)[LSB]	logwtmp(GLIBC_2.0)[LSB]
login(GLIBC_2.0)[LSB]	logout(GLIBC_2.0)[LSB]	openpty(GLIBC_2.0)[LSB]

## A.10 libz

43

The behavior of the interfaces in this library is specified by the following Standards.

44

This Specification [LSB]

45

**Table A-13 libz Function Interfaces**

adler32[LSB]	gzclose[LSB]	gztell[LSB]
compress[LSB]	gzdopen[LSB]	gzwrite[LSB]

compress2[LSB]	gzeof[LSB]	inflate[LSB]
compressBound[LSB]	gzerror[LSB]	inflateEnd[LSB]
crc32[LSB]	gzflush[LSB]	inflateInit2_[LSB]
deflate[LSB]	gzgetc[LSB]	inflateInit_[LSB]
deflateBound[LSB]	gzgets[LSB]	inflateReset[LSB]
deflateCopy[LSB]	gzopen[LSB]	inflateSetDictionary[LSB]
deflateEnd[LSB]	gzprintf[LSB]	inflateSync[LSB]
deflateInit2_[LSB]	gzputc[LSB]	inflateSyncPoint[LSB]
deflateInit_[LSB]	gzputs[LSB]	uncompress[LSB]
deflateParams[LSB]	gzread[LSB]	zError[LSB]
deflateReset[LSB]	gzrewind[LSB]	zlibVersion[LSB]
deflateSetDictionary[LSB] ]	gzseek[LSB]	
get_crc_table[LSB]	gzsetparams[LSB]	

## **Annex B Future Directions (Informative)**

### **B.1 Introduction**

1 This appendix describes interfaces that are under development and aimed at future  
2 releases of this specification. At this stage, such interfaces are at best recommended  
3 practice, and do not constitute normative requirements of this specification.  
4 Applications may not assume that any system provides these interfaces.

5 We encourage system implementors and ISVs to provide these interfaces, and to  
6 provide feedback on their specification to lsbspec@freestandards.org  
7 (mailto://lsb-spec@freestandards.org). These interfaces may well be further  
8 modified during the development process, and may be withdrawn if concensus  
9 cannot be reached.

## B.2 Commands And Utilities

### **lsbinstall**

#### **Name**

10           **lsbinstall** – installation tool for various types of data

#### **Synopsis**

11       `/usr/lib/lsb/lsbinstall [-c | --check | -r | --remove] { -t type | --type=type }`  
 12       `[-p package | --package=package] operand...`

#### **Description**

13       The **lsbinstall** utility may be used to install certain types of files into system specific  
 14       locations, repositories, or databases. This command may be used during a package  
 15       post installation script to add package specific data to system wide repositories. A  
 16       user may need appropriate privilege to invoke **lsbinstall**.

17       The operand (or operands) name an object of type *type* (see below) that belongs to a  
 18       package named *package*. The combination of package name, object type and object  
 19       name should be unique amongst all objects installed by **lsbinstall**. The **lsbinstall**  
 20       utility may rename an object if another package already owns an object of the same  
 21       type with the same name.

22       **Note:** If a namespace collision is detected by **lsbinstall**, it is unspecified how the object is  
 23       renamed, although typical implementations may prepend the package name to the object  
 24       in some way (e.g. *package.obj-name*). The **lsbinstall** utility may maintain a database of  
 25       the mappings it has performed during installation in order to ensure that the correct  
 26       object is removed during a subsequent removal operation.

27       Scripts installed by **lsbinstall** should not make use of the script name in order to  
 28       decide on their functionality.

29       **Note:** It is appropriate for such a script to use the script name in error messages, usage  
 30       statements, etc. The only guarantee made by **lsbinstall** is the effect that an installation (or  
 31       removal) should have, not where a script is installed, or how it is named.

32       The `-p pkg` or `--package=pkg` is required for all object types unless explicitly noted  
 33       below.

34       If the `-c` or `--check` option is specified, **lsbinstall** should test to see if there is an  
 35       existing object of the type specified already installed. If there is, **lsbinstall** should  
 36       print a message to its standard output and immediately exit with a status of zero. If  
 37       there is no object of the type and name specified already installed, **lsbinstall** should  
 38       exit with a non-zero status and take no further action.

39       If the `-r` or `--remove` is specified, the named object of the specified type should be  
 40       removed or disabled from the system, except as noted below. The behavior is  
 41       unspecified if the named object was not previously installed by **lsbinstall**.

42       **Note:** **lsbinstall** may rename objects during installation in order to prevent name  
 43       collisions where another package has already installed an object with the given name.  
 44       Using **lsbinstall --remove** will remove only the object belonging to the named package,  
 45       and not the object belonging to another package.

46                   Also note that the intent of the `--remove` option is to prevent the effect of the installed  
47                   object; it should be sufficient to disable or comment out the addition in some way, while  
48                   leaving the content behind. It is not intended that `--remove` be required to be the exact  
49                   reverse of installation.

## Object Types

50                   The `-t type` or `--type=type` option should support at least the following types:

51                   profile

52                   install a profile script into a system specific location. There should be one  
53                   operand, that names a profile shell script. The behavior is unspecified if this  
54                   name does not have the suffix `.sh`.

55                   The `sh` utility should read and execute commands in its current execution  
56                   environment from all such installed profile shell scripts when invoked as an  
57                   interactive login shell, or if the `-l` (the letter *ell*) is specified (see Shell  
58                   Invocation).

59                   service

60                   ensure a service name and number pair is known to the system service database.  
61                   When installing, there must be at least two operands. The first operand should  
62                   have the format `%d/%s` with the port number and protocol values (e.g. `22/tcp`),  
63                   and the second operand should be the name of the service. Any subsequent  
64                   operands provide aliases for this service. The `-p pkg` or `--package=pkg` option  
65                   is not required for service objects, and is ignored if specified. If any of the `-r`,  
66                   `--remove`, `-c` or `--check` options are specified, there should be a single operand  
67                   identifying the port and protocol values (with the same format as above).

68                   It should not be an error to attempt to add a service name to the system service  
69                   database if that service name already exists for the same port and protocol  
70                   combination. If the port and protocol combination was already present, but the  
71                   name unknown, the name should be added as an alias to the existing entry. It  
72                   should be an error to attempt to add a second entry for a given service name  
73                   and protocol, but where the port number differs from an existing entry.

74                   If the `-r` or `--remove` is specified, the system service database need not be  
75                   updated to remove or disable the named service.

76                   inet

77                   add an entry to the system's network super daemon configuration. If none of  
78                   the `-r`, `--remove`, `-c` or `--check` options are specified, the first operand should  
79                   have the format:

80                   "`%s:%s:%s:%s:%s:%s`"

81                   Otherwise, the first operand should have the format

82                   "`%s:%s`"

83                   The fields in the first operand have the following meaning, in order:

84                   `svc_name`

85                   The name of this service. If the name does not contain a `/`, this should  
86                   match the name of an already installed `service` (see also  
87                   `getservbyname()`). If the name contains a `/` character, the behavior is  
88                   unspecified.

89           **Rationale:** This version of the LSB does not specify `getrpcbyname()` nor the  
 90           existence or format of the `/etc/rpc` file. Therefore, installation of RPC based  
 91           services is not specified at this point. A future version of this specification may  
 92           require names containing a / character to be Remote Procedure Call based  
 93           services.

94           **protocol**  
 95           The name of a protocol. The name should be one of those listed in  
 96           `/etc/protocols`. If this attribute is not specified (i.e. a null value is passed),  
 97           the system should use an implementation defined default protocol.

98           **socket\_type**  
 99           One of the following values:  
 100           **stream**  
 101           the service will use a stream type socket.  
 102           **dgram**  
 103           the service will use a datagram type socket.  
 104           **seqpacket**  
 105           the service will use a sequenced packet type socket.  
 106           This field is not required for the `-c`, `--check`, `-r`, or `--remove` options.

107           **wait\_flag**  
 108           If the value of this attribute is `wait`, once the service is started, no further  
 109           requests for that service will be handled until the service exits. If the value  
 110           is `nowait`, the network super daemon should continue to handle further  
 111           requests for the given service while that service is running.

112           **Note:** If the service has the `socket_type` attribute set to `dgram`, the `wait_flag`  
 113           attribute should be set to `wait`, since such services do not have any distinction  
 114           between the socket used for listening and that used for accepting.

115           This field is not required for the `-c`, `--check`, `-r`, or `--remove` options.

116           **user[ .group ]**  
 117           The name of a user from the user login database, optionally followed by the  
 118           name of a group from the group database. The service started to handle this  
 119           request should run with the privileges of the specified user and group. This  
 120           field is not required for the `-c`, `--check`, `-r`, or `--remove` options.

121           **server [ arg ... ]**  
 122           The name of a program to run to handle the request, optionally followed by  
 123           any arguments required. The server name and each of its arguments is  
 124           separated by whitespace. This field is not required for the `-c`, `--check`, `-r`,  
 125           or `--remove` options.

126           If the implementation supports additional controls over services started  
 127           through the inet super daemon, there may be additional,  
 128           implementation-defined, operands.

129                   **Rationale:** Systems that use the **xinetd** super daemon may support additional controls  
130                   such as IP address restrictions, logging requirements, etc. The LSB does not require  
131                   these additional controls. However, it was believed to be of sufficient benefit that  
132                   implementations are granted permission to extend this interface as required.

## Examples

133       lsbinstall --package=myapp --type=profile myco.com-prod.sh  
134       Install the profile shell script for myco.com-prod.sh, part of the myapp package..  
135       lsbinstall --package=myapp --check --type=profile myco.com-prod.sh  
136       Test to see if the profile shell script for myco.com-prod.sh, as part of the myapp  
137       package, is installed correctly.

## Exit Status

138       If the *-c* or *--check* option is specified, **lsbinstall** should exit with a zero status if an  
139       object of the specified type and name is already installed, or non-zero otherwise.  
140       Otherwise, **lsbinstall** should exit with a zero status if the object with the specified  
141       type and name was successfully installed (or removed if the *-r* or *--remove* option  
142       was specified), and non-zero if the installation (or removal) failed. On failure, a  
143       diagnostic message should be printed to the standard error file descriptor.

## Annex C GNU Free Documentation License (Informative)

1 This specification is published under the terms of the GNU Free Documentation  
2 License, Version 1.1, March 2000

3 Copyright (C) 2000 Free Software Foundation, Inc. 59 Temple Place, Suite 330, Boston,  
4 MA 02111-1307 USA Everyone is permitted to copy and distribute verbatim copies of  
5 this license document, but changing it is not allowed.

### C.1 PREAMBLE

6 The purpose of this License is to make a manual, textbook, or other written  
7 document "free" in the sense of freedom: to assure everyone the effective freedom to  
8 copy and redistribute it, with or without modifying it, either commercially or  
9 noncommercially. Secondarily, this License preserves for the author and publisher a  
10 way to get credit for their work, while not being considered responsible for  
11 modifications made by others.

12 This License is a kind of "copyleft", which means that derivative works of the  
13 document must themselves be free in the same sense. It complements the GNU  
14 General Public License, which is a copyleft license designed for free software.

15 We have designed this License in order to use it for manuals for free software,  
16 because free software needs free documentation: a free program should come with  
17 manuals providing the same freedoms that the software does. But this License is not  
18 limited to software manuals; it can be used for any textual work, regardless of  
19 subject matter or whether it is published as a printed book. We recommend this  
20 License principally for works whose purpose is instruction or reference.

### C.2 APPLICABILITY AND DEFINITIONS

21 This License applies to any manual or other work that contains a notice placed by  
22 the copyright holder saying it can be distributed under the terms of this License. The  
23 "Document", below, refers to any such manual or work. Any member of the public is  
24 a licensee, and is addressed as "you".

25 A "Modified Version" of the Document means any work containing the Document or  
26 a portion of it, either copied verbatim, or with modifications and/or translated into  
27 another language.

28 A "Secondary Section" is a named appendix or a front-matter section of the  
29 Document that deals exclusively with the relationship of the publishers or authors of  
30 the Document to the Document's overall subject (or to related matters) and contains  
31 nothing that could fall directly within that overall subject. (For example, if the  
32 Document is in part a textbook of mathematics, a Secondary Section may not explain  
33 any mathematics.) The relationship could be a matter of historical connection with  
34 the subject or with related matters, or of legal, commercial, philosophical, ethical or  
35 political position regarding them.

36 The "Invariant Sections" are certain Secondary Sections whose titles are designated,  
37 as being those of Invariant Sections, in the notice that says that the Document is  
38 released under this License.

39 The "Cover Texts" are certain short passages of text that are listed, as Front-Cover  
40 Texts or Back-Cover Texts, in the notice that says that the Document is released  
41 under this License.

42 A "Transparent" copy of the Document means a machine-readable copy, represented  
43 in a format whose specification is available to the general public, whose contents can  
44 be viewed and edited directly and straightforwardly with generic text editors or (for  
45 images composed of pixels) generic paint programs or (for drawings) some widely  
46 available drawing editor, and that is suitable for input to text formatters or for  
47 automatic translation to a variety of formats suitable for input to text formatters. A  
48 copy made in an otherwise Transparent file format whose markup has been  
49 designed to thwart or discourage subsequent modification by readers is not  
50 Transparent. A copy that is not "Transparent" is called "Opaque".

51 Examples of suitable formats for Transparent copies include plain ASCII without  
52 markup, Texinfo input format, LaTeX input format, SGML or XML using a publicly  
53 available DTD, and standard-conforming simple HTML designed for human  
54 modification. Opaque formats include PostScript, PDF, proprietary formats that can  
55 be read and edited only by proprietary word processors, SGML or XML for which  
56 the DTD and/or processing tools are not generally available, and the  
57 machine-generated HTML produced by some word processors for output purposes  
58 only.

59 The "Title Page" means, for a printed book, the title page itself, plus such following  
60 pages as are needed to hold, legibly, the material this License requires to appear in  
61 the title page. For works in formats which do not have any title page as such, "Title  
62 Page" means the text near the most prominent appearance of the work's title,  
63 preceding the beginning of the body of the text.

### C.3 VERBATIM COPYING

64 You may copy and distribute the Document in any medium, either commercially or  
65 noncommercially, provided that this License, the copyright notices, and the license  
66 notice saying this License applies to the Document are reproduced in all copies, and  
67 that you add no other conditions whatsoever to those of this License. You may not  
68 use technical measures to obstruct or control the reading or further copying of the  
69 copies you make or distribute. However, you may accept compensation in exchange  
70 for copies. If you distribute a large enough number of copies you must also follow  
71 the conditions in section 3.

72 You may also lend copies, under the same conditions stated above, and you may  
73 publicly display copies.

### C.4 COPYING IN QUANTITY

74 If you publish printed copies of the Document numbering more than 100, and the  
75 Document's license notice requires Cover Texts, you must enclose the copies in  
76 covers that carry, clearly and legibly, all these Cover Texts: Front-Cover Texts on the  
77 front cover, and Back-Cover Texts on the back cover. Both covers must also clearly  
78 and legibly identify you as the publisher of these copies. The front cover must  
79 present the full title with all words of the title equally prominent and visible. You  
80 may add other material on the covers in addition. Copying with changes limited to  
81 the covers, as long as they preserve the title of the Document and satisfy these  
82 conditions, can be treated as verbatim copying in other respects.

83 If the required texts for either cover are too voluminous to fit legibly, you should put  
84 the first ones listed (as many as fit reasonably) on the actual cover, and continue the  
85 rest onto adjacent pages.

86 If you publish or distribute Opaque copies of the Document numbering more than  
87 100, you must either include a machine-readable Transparent copy along with each

88 Opaque copy, or state in or with each Opaque copy a publicly-accessible  
 89 computer-network location containing a complete Transparent copy of the  
 90 Document, free of added material, which the general network-using public has  
 91 access to download anonymously at no charge using public-standard network  
 92 protocols. If you use the latter option, you must take reasonably prudent steps, when  
 93 you begin distribution of Opaque copies in quantity, to ensure that this Transparent  
 94 copy will remain thus accessible at the stated location until at least one year after the  
 95 last time you distribute an Opaque copy (directly or through your agents or  
 96 retailers) of that edition to the public.

97 It is requested, but not required, that you contact the authors of the Document well  
 98 before redistributing any large number of copies, to give them a chance to provide  
 99 you with an updated version of the Document.

## C.5 MODIFICATIONS

100 You may copy and distribute a Modified Version of the Document under the  
 101 conditions of sections 2 and 3 above, provided that you release the Modified Version  
 102 under precisely this License, with the Modified Version filling the role of the  
 103 Document, thus licensing distribution and modification of the Modified Version to  
 104 whoever possesses a copy of it. In addition, you must do these things in the  
 105 Modified Version:

- 106 A. Use in the Title Page (and on the covers, if any) a title distinct from that of the  
     107 Document, and from those of previous versions (which should, if there were  
     108 any, be listed in the History section of the Document). You may use the same  
     109 title as a previous version if the original publisher of that version gives  
     110 permission.
- 111 B. List on the Title Page, as authors, one or more persons or entities responsible  
     112 for authorship of the modifications in the Modified Version, together with at  
     113 least five of the principal authors of the Document (all of its principal authors,  
     114 if it has less than five).
- 115 C. State on the Title page the name of the publisher of the Modified Version, as  
     116 the publisher.
- 117 D. Preserve all the copyright notices of the Document.
- 118 E. Add an appropriate copyright notice for your modifications adjacent to the  
     119 other copyright notices.
- 120 F. Include, immediately after the copyright notices, a license notice giving the  
     121 public permission to use the Modified Version under the terms of this License,  
     122 in the form shown in the Addendum below.
- 123 G. Preserve in that license notice the full lists of Invariant Sections and required  
     124 Cover Texts given in the Document's license notice.
- 125 H. Include an unaltered copy of this License.
- 126 I. Preserve the section entitled "History", and its title, and add to it an item  
     127 stating at least the title, year, new authors, and publisher of the Modified  
     128 Version as given on the Title Page. If there is no section entitled "History" in  
     129 the Document, create one stating the title, year, authors, and publisher of the  
     130 Document as given on its Title Page, then add an item describing the Modified  
     131 Version as stated in the previous sentence.
- 132 J. Preserve the network location, if any, given in the Document for public access  
     133 to a Transparent copy of the Document, and likewise the network locations

- 134 given in the Document for previous versions it was based on. These may be  
135 placed in the "History" section. You may omit a network location for a work  
136 that was published at least four years before the Document itself, or if the  
137 original publisher of the version it refers to gives permission.
- 138 K. In any section entitled "Acknowledgements" or "Dedications", preserve the  
139 section's title, and preserve in the section all the substance and tone of each of  
140 the contributor acknowledgements and/or dedications given therein.
- 141 L. Preserve all the Invariant Sections of the Document, unaltered in their text and  
142 in their titles. Section numbers or the equivalent are not considered part of the  
143 section titles.
- 144 M. Delete any section entitled "Endorsements". Such a section may not be  
145 included in the Modified Version.
- 146 N. Do not retitle any existing section as "Endorsements" or to conflict in title with  
147 any Invariant Section.
- 148 If the Modified Version includes new front-matter sections or appendices that  
149 qualify as Secondary Sections and contain no material copied from the Document,  
150 you may at your option designate some or all of these sections as invariant. To do  
151 this, add their titles to the list of Invariant Sections in the Modified Version's license  
152 notice. These titles must be distinct from any other section titles.
- 153 You may add a section entitled "Endorsements", provided it contains nothing but  
154 endorsements of your Modified Version by various parties—for example, statements  
155 of peer review or that the text has been approved by an organization as the  
156 authoritative definition of a standard.
- 157 You may add a passage of up to five words as a Front-Cover Text, and a passage of  
158 up to 25 words as a Back-Cover Text, to the end of the list of Cover Texts in the  
159 Modified Version. Only one passage of Front-Cover Text and one of Back-Cover  
160 Text may be added by (or through arrangements made by) any one entity. If the  
161 Document already includes a cover text for the same cover, previously added by you  
162 or by arrangement made by the same entity you are acting on behalf of, you may not  
163 add another; but you may replace the old one, on explicit permission from the  
164 previous publisher that added the old one.
- 165 The author(s) and publisher(s) of the Document do not by this License give  
166 permission to use their names for publicity for or to assert or imply endorsement of  
167 any Modified Version.

## C.6 COMBINING DOCUMENTS

- 168 You may combine the Document with other documents released under this License,  
169 under the terms defined in section 4 above for modified versions, provided that you  
170 include in the combination all of the Invariant Sections of all of the original  
171 documents, unmodified, and list them all as Invariant Sections of your combined  
172 work in its license notice.
- 173 The combined work need only contain one copy of this License, and multiple  
174 identical Invariant Sections may be replaced with a single copy. If there are multiple  
175 Invariant Sections with the same name but different contents, make the title of each  
176 such section unique by adding at the end of it, in parentheses, the name of the  
177 original author or publisher of that section if known, or else a unique number. Make  
178 the same adjustment to the section titles in the list of Invariant Sections in the license  
179 notice of the combined work.

180 In the combination, you must combine any sections entitled "History" in the various  
 181 original documents, forming one section entitled "History"; likewise combine any  
 182 sections entitled "Acknowledgements", and any sections entitled "Dedications". You  
 183 must delete all sections entitled "Endorsements."

## C.7 COLLECTIONS OF DOCUMENTS

184 You may make a collection consisting of the Document and other documents  
 185 released under this License, and replace the individual copies of this License in the  
 186 various documents with a single copy that is included in the collection, provided  
 187 that you follow the rules of this License for verbatim copying of each of the  
 188 documents in all other respects.

189 You may extract a single document from such a collection, and distribute it  
 190 individually under this License, provided you insert a copy of this License into the  
 191 extracted document, and follow this License in all other respects regarding verbatim  
 192 copying of that document.

## C.8 AGGREGATION WITH INDEPENDENT WORKS

193 A compilation of the Document or its derivatives with other separate and  
 194 independent documents or works, in or on a volume of a storage or distribution  
 195 medium, does not as a whole count as a Modified Version of the Document,  
 196 provided no compilation copyright is claimed for the compilation. Such a  
 197 compilation is called an "aggregate", and this License does not apply to the other  
 198 self-contained works thus compiled with the Document, on account of their being  
 199 thus compiled, if they are not themselves derivative works of the Document.

200 If the Cover Text requirement of section 3 is applicable to these copies of the  
 201 Document, then if the Document is less than one quarter of the entire aggregate, the  
 202 Document's Cover Texts may be placed on covers that surround only the Document  
 203 within the aggregate. Otherwise they must appear on covers around the whole  
 204 aggregate.

## C.9 TRANSLATION

205 Translation is considered a kind of modification, so you may distribute translations  
 206 of the Document under the terms of section 4. Replacing Invariant Sections with  
 207 translations requires special permission from their copyright holders, but you may  
 208 include translations of some or all Invariant Sections in addition to the original  
 209 versions of these Invariant Sections. You may include a translation of this License  
 210 provided that you also include the original English version of this License. In case of  
 211 a disagreement between the translation and the original English version of this  
 212 License, the original English version will prevail.

## C.10 TERMINATION

213 You may not copy, modify, sublicense, or distribute the Document except as  
 214 expressly provided for under this License. Any other attempt to copy, modify,  
 215 sublicense or distribute the Document is void, and will automatically terminate your  
 216 rights under this License. However, parties who have received copies, or rights,  
 217 from you under this License will not have their licenses terminated so long as such  
 218 parties remain in full compliance.

## C.11 FUTURE REVISIONS OF THIS LICENSE

219 The Free Software Foundation may publish new, revised versions of the GNU Free  
220 Documentation License from time to time. Such new versions will be similar in spirit  
221 to the present version, but may differ in detail to address new problems or concerns.  
222 See <http://www.gnu.org/copyleft/>.

223 Each version of the License is given a distinguishing version number. If the  
224 Document specifies that a particular numbered version of this License "or any later  
225 version" applies to it, you have the option of following the terms and conditions  
226 either of that specified version or of any later version that has been published (not as  
227 a draft) by the Free Software Foundation. If the Document does not specify a version  
228 number of this License, you may choose any version ever published (not as a draft)  
229 by the Free Software Foundation.

## C.12 How to use this License for your documents

230 To use this License in a document you have written, include a copy of the License in  
231 the document and put the following copyright and license notices just after the title  
232 page:

233 Copyright (c) YEAR YOUR NAME. Permission is granted to copy, distribute and/or  
234 modify this document under the terms of the GNU Free Documentation License, Version  
235 1.1 or any later version published by the Free Software Foundation; with the Invariant  
236 Sections being LIST THEIR TITLES, with the Front-Cover Texts being LIST, and with the  
237 Back-Cover Texts being LIST. A copy of the license is included in the section entitled  
238 "GNU Free Documentation License".

239 If you have no Invariant Sections, write "with no Invariant Sections" instead of  
240 saying which ones are invariant. If you have no Front-Cover Texts, write "no  
241 Front-Cover Texts" instead of "Front-Cover Texts being LIST"; likewise for  
242 Back-Cover Texts.

243 If your document contains nontrivial examples of program code, we recommend  
244 releasing these examples in parallel under your choice of free software license, such  
245 as the GNU General Public License, to permit their use in free software.