

The Open Group Standard

Additional APIs for the Base Specifications Issue 8, Part 2

THE *Open* GROUP

Copyright © 2022, The Open Group

The Open Group hereby authorizes you to use this document for any purpose, PROVIDED THAT any copy of this document, or any part thereof, which you make shall retain all copyright and other proprietary notices contained herein.

This document may contain other proprietary notices and copyright information.

Nothing contained herein shall be construed as conferring by implication, estoppel, or otherwise any license or right under any patent or trademark of The Open Group or any third party. Except as expressly provided above, nothing contained herein shall be construed as conferring any license or right under any copyright of The Open Group.

Note that any product, process, or technology in this document may be the subject of other intellectual property rights reserved by The Open Group, and may not be licensed hereunder.

This document is provided “AS IS” WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, OR NON-INFRINGEMENT. Some jurisdictions do not allow the exclusion of implied warranties, so the above exclusion may not apply to you.

Any publication of The Open Group may include technical inaccuracies or typographical errors. Changes may be periodically made to these publications; these changes will be incorporated in new editions of these publications. The Open Group may make improvements and/or changes in the products and/or the programs described in these publications at any time without notice.

Should any viewer of this document respond with information including feedback data, such as questions, comments, suggestions, or the like regarding the content of this document, such information shall be deemed to be non-confidential and The Open Group shall have no obligation of any kind with respect to such information and shall be free to reproduce, use, disclose, and distribute the information to others without limitation. Further, The Open Group shall be free to use any ideas, concepts, know-how, or techniques contained in such information for any purpose whatsoever including but not limited to developing, manufacturing, and marketing products incorporating such information.

If you did not obtain this copy through The Open Group, it may not be the latest version. For your convenience, the latest version of this publication may be downloaded at www.opengroup.org/library.

The Open Group Standard

Additional APIs for the Base Specifications Issue 8, Part 2

ISBN: TBA

Document Number: C22x

Published by The Open Group, <Month> 2022

Comments relating to the material contained in this document may be submitted to:

The Open Group, Apex Plaza, Forbury Road, Reading, Berkshire, RG1 1AX, United Kingdom

or by electronic mail to:

ogspecs@opengroup.org

Contents

1	Introduction.....	1
1.1	Scope.....	1
1.2	Relationship to Other Formal Standards.....	1
2	Application Program Interfaces	2
2.1	Change Bars.....	2
2.2	Reference Pages.....	2

DRAFT

Preface

The Open Group

The Open Group is a global consortium that enables the achievement of business objectives through technology standards. Our diverse membership of more than 870 organizations includes customers, systems and solutions suppliers, tools vendors, integrators, academics, and consultants across multiple industries.

The mission of The Open Group is to drive the creation of Boundaryless Information Flow™ achieved by:

- Working with customers to capture, understand, and address current and emerging requirements, establish policies, and share best practices
- Working with suppliers, consortia, and standards bodies to develop consensus and facilitate interoperability, to evolve and integrate specifications and open source technologies
- Offering a comprehensive set of services to enhance the operational efficiency of consortia
- Developing and operating the industry's premier certification service and encouraging procurement of certified products

Further information on The Open Group is available at www.opengroup.org.

The Open Group publishes a wide range of technical documentation, most of which is focused on development of Standards and Guides, but which also includes white papers, technical studies, certification and testing documentation, and business titles. Full details and a catalog are available at www.opengroup.org/library.

This Document

This document has been prepared by The Open Group Base Working Group. The Open Group Base Working Group is considering submitting a number of additional APIs to the Austin Group as input to the Issue 8 revision of the Base Specifications.

This document contains the second set of these APIs.

Trademarks

ArchiMate, DirecNet, Making Standards Work, Open O logo, Open O and Check Certification logo, Platform 3.0, The Open Group, TOGAF, UNIX, UNIXWARE, and the Open Brand X logo are registered trademarks and Boundaryless Information Flow, Build with Integrity Buy with Confidence, Commercial Aviation Reference Architecture, Dependability Through Assuredness, Digital Practitioner Body of Knowledge, DPBoK, EMMM, FACE, the FACE logo, FHIM Profile Builder, the FHIM logo, FPB, Future Airborne Capability Environment, IT4IT, the IT4IT logo, O-AA, O-DEF, O-HERA, O-PAS, Open Agile Architecture, Open FAIR, Open Footprint, Open Process Automation, Open Subsurface Data Universe, Open Trusted Technology Provider, OSDU, Sensor Integration Simplified, SOSA, and the SOSA logo are trademarks of The Open Group.

All other brands, company, and product names are used for identification purposes only and may be trademarks that are the sole property of their respective owners.

Acknowledgements

The Open Group gratefully acknowledges the contribution of the following in the development of this document:

- The Open Group Base Working Group
- The Austin Group

The Open Group gratefully acknowledges the following reviewers who participated in the Company Review of this document:

- TBA

DRAFT

1 Introduction

1.1 Scope

The purpose of this document is to define a set of additional APIs for inclusion in the Issue 8 revision of the Base Specifications of the Single UNIX[®] Specification.

The additional APIs proposed by participants in the Austin Group that The Open Group has agreed to sponsor are as follows:

Header

`<libintl.h>`

Functions

<i>bind_textdomain_codeset()</i>	<i>getresgid()</i>
<i>bindtextdomain()</i>	<i>getresuid()</i>
<i>dcgettext()</i>	<i>gettext()</i>
<i>dcgettext_l()</i>	<i>gettext_l()</i>
<i>dcngettext()</i>	<i>ngettext()</i>
<i>dcngettext_l()</i>	<i>ngettext_l()</i>
<i>dgettext()</i>	<i>setresgid()</i>
<i>dgettext_l()</i>	<i>setresuid()</i>
<i>dngettext()</i>	<i>textdomain()</i>
<i>dngettext_l()</i>	

Utilities

<i>gettext</i>	<i>realpath</i>
<i>msgfmt</i>	<i>timeout</i>
<i>ngettext</i>	<i>xgettext</i>
<i>readlink</i>	

1.2 Relationship to Other Formal Standards

This Standard is being forwarded to the Austin Group for consideration as input to the Issue 8 revision of the Base Specifications.

2 Application Program Interfaces

The following pages are extracted from a complete draft of the Base Specifications in which the proposed changes have been applied, with change bars showing the differences from Issue 8 Draft 2.1. Only pages with technical changes are included – editorial changes such as additions to SEE ALSO and CHANGE HISTORY sections have been omitted (unless they appear on the same page as a technical change). The complete draft is also being made available for reference.

As a consequence of the change to NLSPATH in XBD Section 8.2, a change will also need to be made to the NLSPATH description on all existing utility reference pages. These changes are not included here but will be made during the preparation of Issue 8 Draft 3.

2.1 Change Bars

Changed lines are marked with a '|' in the right-hand margin, new lines with a '+', and deleted lines with a '-'.

Note that sometimes the placement of change bars is slightly inaccurate. In particular, changes may extend into a line following a set of change-barred lines. Also, changes within tables do not have change bars.

2.2 Reference Pages

The reference pages for the new header, function, and utility additions, and pages with related changes follow.

1428 3.106 Dot

1429 In the context of naming files, the filename consisting of a single <period> character (' . ').

1430 **Note:** In the context of shell special built-in utilities, see *dot* in XCU [Section 2.14](#) (on page 2382).

1431 Pathname Resolution is defined in detail in [Section 4.14](#) (on page 93).

1432 3.107 Dot-Dot

1433 The filename consisting solely of two <period> characters (" . . ").

1434 **Note:** Pathname Resolution is defined in detail in [Section 4.14](#) (on page 93).

1435 3.108 Dot-Po File

1436 See *Portable Messages Object Source File* in [Section 3.257](#) (on page 65).

1437 3.109 Double-Quote Character

1438 The character ' " ', also known as <quotation-mark>.

1439 **Note:** The “double” adjective in this term refers to the two strokes in the character glyph.
1440 POSIX.1-202x never uses the term “double-quote” to refer to two apostrophes or quotation-
1441 marks.

1442 3.110 Downshifting

1443 The conversion of an uppercase character that has a single-character lowercase representation
1444 into this lowercase representation.

1445 3.111 Driver

1446 A module that controls data transferred to and received from devices.

1447 **Note:** Drivers are traditionally written to be a part of the system implementation, although they are
1448 frequently written separately from the writing of the implementation. A driver may contain
1449 processor-specific code, and therefore be non-portable.

1450 3.112 Effective Group ID

1451 An attribute of a process that is used in determining various permissions, including file access
1452 permissions; see also [Section 3.161](#) (on page 51).

1453 3.113 Effective User ID

1454 An attribute of a process that is used in determining various permissions, including file access
1455 permissions; see also [Section 3.394](#) (on page 84).

1786 3.197 Message Catalog

1787 In the context of providing natural language messages to the user, a file or storage area
1788 containing program messages, command prompts, and responses to prompts for a particular
1789 native language, territory, and codeset.

1790 3.198 Message Catalog Descriptor

1791 In the context of providing natural language messages to the user, a per-process unique value
1792 used to identify an open message catalog. A message catalog descriptor may be implemented
1793 using a file descriptor.

1794 3.199 Message Queue

1795 In the context of programmatic message passing, an object to which messages can be added and
1796 removed. Messages may be removed in the order in which they were added or in priority order.

1797 3.200 Messages Object +

1798 A file containing message identifiers and translations in an unspecified format. Used by the *+*
1799 *gettext* family of functions and the *gettext* and *ngettext* utilities for internationalization and *+*
1800 localization of programs and scripts. Messages objects have the filename suffix *.mo*, and can be *+*
1801 created by the *msgfmt* utility. *+*

1802 See also [Section 3.374](#) (on page 81).

1803 3.201 Mode

1804 A collection of attributes that specifies a file's type and its access permissions.

1805 **Note:** File Access Permissions are defined in detail in [Section 4.6](#) (on page 90).

1806 3.202 Monotonic Clock

1807 A clock measuring real time, whose value cannot be set via *clock_settime()* and which cannot
1808 have negative clock jumps.

1809 3.203 Mount Point

1810 Either the system root directory or a directory for which the *st_dev* field of structure **stat** differs
1811 from that of its parent directory.

1812 **Note:** The **stat** structure is defined in detail in [<sys/stat.h>](#).

1813 3.204 Multi-Character Collating Element

1814 A sequence of two or more characters that collate as an entity. For example, in some coded

2023 **3.252 Pipe**

2024 An object identical to a FIFO which has no links in the file hierarchy.

2025 **Note:** The *pipe()* function is defined in detail in the System Interfaces volume of POSIX.1-202x.2026 **3.253 Polling**2027 A scheduling scheme whereby the local process periodically checks until the pre-specified
2028 events (for example, read, write) have occurred.2029 **3.254 Portable Character Set**2030 The collection of characters that are required to be present in all locales supported by
2031 conforming systems.2032 **Note:** The Portable Character Set is defined in detail in [Section 6.1](#) (on page 105).2033 This term is contrasted against the smaller portable filename character set; see also [Section 3.256](#).2034 **3.255 Portable Filename**

2035 A filename consisting only of characters from the portable filename character set.

2036 **Note:** Applications should avoid using filenames that have the <hyphen-minus> character as the first
2037 character since this may cause problems when filenames are passed as command line
2038 arguments.2039 **3.256 Portable Filename Character Set**

2040 The set of characters from which portable filenames are constructed.

2041 A B C D E F G H I J K L M N O P Q R S T U V W X Y Z
2042 a b c d e f g h i j k l m n o p q r s t u v w x y z
2043 0 1 2 3 4 5 6 7 8 9 . _ -2044 The last three characters are the <period>, <underscore>, and <hyphen-minus> characters,
2045 respectively. See also [Section 3.245](#) (on page 63).2046 **3.257 Portable Messages Object Source File (or Dot-Po File)** +2047 A text file containing messages and directives. A portable messages object source file can be +
2048 compiled into a messages object by the *msgfmt* utility. +2049 **Note:** By convention, portable messages object source files have filenames ending with the **.po** suffix. +
2050 Utility descriptions in this standard frequently use dot-po file as a shorthand for portable +
2051 messages object source file (even though the **.po** suffix need not be included in the filename). +
2052 Template portable messages object source files can be created from C-language source files by +
2053 the *xgettext* utility. +

2477 3.368 System Process

2478 An object other than a process executing an application, that is provided by the system and has a
2479 process ID.

2480 3.369 System Reboot

2481 See *System Boot* defined in [Section 3.362](#) (on page 80).

2482 3.370 System-Wide

2483 Pertaining to events occurring in all processes existing in an implementation at a given point in
2484 time.

2485 3.371 Tab Character (<tab>)

2486 A character that in the output stream indicates that printing or displaying should start at the
2487 next horizontal tabulation position on the current line. It is the character designated by '\t' in
2488 the C language. If the current position is at or past the last defined horizontal tabulation
2489 position, the behavior is unspecified. It is unspecified whether this character is the exact
2490 sequence transmitted to an output device by the system to accomplish the tabulation.

2491 3.372 Terminal (or Terminal Device)

2492 A character special file that obeys the specifications of the general terminal interface.

2493 **Note:** The General Terminal Interface is defined in detail in [Chapter 11](#) (on page 185).

2494 3.373 Text Column

2495 A roughly rectangular block of characters capable of being laid out side-by-side next to other
2496 text columns on an output page or terminal screen. The widths of text columns are measured in
2497 column positions.

2498 3.374 Text Domain

2499 A named collection of messages objects (one messages object per supported language) for
2500 internationalization and localization purposes. A text domain is often named after the
2501 application or library that provides the collection, but may have a more general name if it is
2502 intended to be shared by multiple applications or libraries. +

2503 **Note:** The use of text domains is defined in detail in the descriptions of the [bindtextdomain\(\)](#) and
2504 [gettext](#) family of functions in the System Interfaces volume of POSIX.1-202x. +

2505 3.375 Text File

2506 A file that contains characters organized into zero or more lines. The lines do not contain NUL

Locale

3609 7.1 General

3610 A locale is the definition of the subset of a user's environment that depends on language and
 3611 cultural conventions. It is made up from one or more categories. Each category is identified by
 3612 its name and controls specific aspects of the behavior of components of the system. Category
 3613 names correspond to the following environment variable names:

3614 *LC_CTYPE* Character classification and case conversion.

3615 *LC_COLLATE* Collation order.

3616 *LC_MONETARY* Monetary formatting.

3617 *LC_NUMERIC* Numeric, non-monetary formatting.

3618 *LC_TIME* Date and time formats.

3619 *LC_MESSAGES* Formats of informative and diagnostic messages and interactive responses.

3620 The standard utilities in the Shell and Utilities volume of POSIX.1-202x shall base their behavior
 3621 on the current locale, as defined in the ENVIRONMENT VARIABLES section for each utility.
 3622 The behavior of some of the C-language functions defined in the System Interfaces volume of
 3623 POSIX.1-202x shall also be modified based on a locale selection. The locale to be used by these
 3624 functions can be selected in the following ways:

- 3625 1. For functions such as *isalnum_l()* that take a locale object as an argument, a locale object
 3626 can be obtained from *newlocale()* or *duplocale()* and passed to the function.
- 3627 2. For functions that do not take a locale object as an argument, the current locale for the
 3628 thread can be set by calling *uselocale()* or the global locale for the process can be set by
 3629 calling *setlocale()*. Such functions shall use the current locale of the calling thread if one
 3630 has been set for that thread; otherwise, they shall use the global locale.
- 3631 3. Some functions, such as *catopen()* and those related to text domains, may reference +
 3632 various environment variables and a locale category of a specific locale to access files they +
 3633 need to use.

3634 Locales other than those supplied by the implementation can be created via the *localedef* utility,
 3635 provided that the *_POSIX2_LOCALEDEF* symbol is defined on the system. Even if *localedef* is
 3636 not provided, all implementations conforming to the System Interfaces volume of POSIX.1-202x
 3637 shall provide one or more locales that behave as described in this chapter. The input to the
 3638 utility is described in [Section 7.3](#) (on page 116). The value that is used to specify a locale when
 3639 using environment variables shall be the string specified as the *name* operand to the *localedef*
 3640 utility when the locale was created. The strings "C" and "POSIX" are reserved as identifiers for
 3641 the POSIX locale (see [Section 7.2](#), on page 116). When the value of a locale environment variable
 3642 begins with a <slash> (' / '), it shall be interpreted as the pathname of the locale definition; the
 3643 type of file (regular, directory, and so on) used to store the locale definition is implementation-
 3644 defined. If the value does not begin with a <slash>, the mechanism used to locate the locale is
 3645 implementation-defined.

5433 8.2 Internationalization Variables

5434 This section describes environment variables that are relevant to the operation of
5435 internationalized interfaces described in POSIX.1-202x.

5436 Users may use the following environment variables to announce specific localization
5437 requirements to applications. Applications can retrieve this information using the *setlocale()*
5438 function to initialize the correct behavior of the internationalized interfaces. The descriptions of
5439 the internationalization environment variables describe the resulting behavior only when the
5440 application locale is initialized in this way. The use of the internationalization variables by
5441 utilities described in the Shell and Utilities volume of POSIX.1-202x is described in the
5442 ENVIRONMENT VARIABLES section for those utilities in addition to the global effects
5443 described in this section.

5444 **LANG** This variable shall determine the locale category for native language, local
5445 customs, and coded character set in the absence of the *LC_ALL* and other *LC_**
5446 (*LC_COLLATE*, *LC_CTYPE*, *LC_MESSAGES*, *LC_MONETARY*, *LC_NUMERIC*,
5447 *LC_TIME*) environment variables. This can be used by applications to determine
5448 the language to use for error messages and instructions, collating sequences, date
5449 formats, and so on.

5450 **LANGUAGE** +

5451 The *LANGUAGE* environment variable shall be examined to determine the +
5452 messages object to be used for the *gettext* family of functions or the *gettext* and +
5453 *gettext* utilities if *NLSPATH* is not set or the evaluation of *NLSPATH* did not lead +
5454 to a suitable messages object being found. The value of *LANGUAGE* shall be a list +
5455 of locale names separated by a <colon> (':') character. If *LANGUAGE* is set to a +
5456 non-empty string, each locale name shall be tried in the specified order and if a +
5457 messages object is found, it shall be used for translation. If a locale name has the +
5458 format *language[_territory][.codeset][@modifier]*, additional searches of locale names +
5459 without *.codeset* (if present), without *_territory* (if present), and without *@modifier* (if +
5460 present) may be performed; if *.codeset* is not present, additional searches of locale +
5461 names with an added *.codeset* may be performed. If locale names contain a <slash> +
5462 ('/') character, or consist entirely of a dot (".") or dot-dot (". .") character +
5463 sequence, or are empty the behavior is implementation defined and they may be +
5464 ignored for security reasons. +

5465 The locale names in *LANGUAGE* shall override the locale name associated with the +
5466 "active category" of the current locale or, in the case of functions with an *_l* +
5467 suffix, the provided locale object, and the language-specific part of the default +
5468 search path for messages objects, unless the locale name that would be overridden +
5469 is C or POSIX. For the *dcgettext()*, *dcgettext_l()*, *dcngettext()*, and *dcngettext_l()* +
5470 functions, the active category is specified by the *category* argument; for all other +
5471 *gettext* family functions and for the *gettext* and *gettext* utilities, the active category +
5472 is *LC_MESSAGES*. +

5473 For example, if: +

- 5474 • The *LC_MESSAGES* environment variable is "de_DE" (and *LC_ALL* is unset) +
- 5475 and *setlocale(LC_ALL, "")* has been used to set the current locale +
- 5476 • The *LANGUAGE* environment variable is "fr_FR:it" +
- 5477 • Messages objects are by default searched for in */gettextlib* +

5478 then the following pathnames are tried in this order by *gettext* family functions that +
5479 have neither a category argument nor an *_l* suffix until a valid messages object is +
5480 found: +

5481		• <code>/gettextlib/fr_FR/LC_MESSAGES/textdomain.mo</code>	+
5482		• (Optionally) <code>/gettextlib/fr/LC_MESSAGES/textdomain.mo</code>	+
5483		• (Optionally) the above two pathnames with added <code>.codeset</code> elements	+
5484		• <code>/gettextlib/it/LC_MESSAGES/textdomain.mo</code>	+
5485		• (Optionally) the above pathname with added <code>.codeset</code> elements	+
5486		• <code>/gettextlib/de_DE/LC_MESSAGES/textdomain.mo</code>	+
5487	<code>LC_ALL</code>	This variable shall determine the values for all locale categories. The value of the <code>LC_ALL</code> environment variable has precedence over any of the other environment variables starting with <code>LC_</code> (<code>LC_COLLATE</code> , <code>LC_CTYPE</code> , <code>LC_MESSAGES</code> , <code>LC_MONETARY</code> , <code>LC_NUMERIC</code> , <code>LC_TIME</code>) and the <code>LANG</code> environment variable.	
5488			
5489			
5490			
5491	<code>LC_COLLATE</code>		
5492		This variable shall determine the locale category for character collation. It determines collation information for regular expressions and sorting, including equivalence classes and multi-character collating elements, in various utilities and the <code>strcoll()</code> and <code>strxfrm()</code> functions. Additional semantics of this variable, if any, are implementation-defined.	
5493			
5494			
5495			
5496			
5497	<code>LC_CTYPE</code>	This variable shall determine the locale category for character handling functions, such as <code>tolower()</code> , <code>toupper()</code> , and <code>isalpha()</code> . This environment variable determines the interpretation of sequences of bytes of text data as characters (for example, single as opposed to multi-byte characters), the classification of characters (for example, alpha, digit, graph), and the behavior of character classes. Additional semantics of this variable, if any, are implementation-defined.	
5498			
5499			
5500			
5501			
5502			
5503	<code>LC_MESSAGES</code>		
5504		This variable shall determine the locale category for processing affirmative and negative responses and the language and cultural conventions in which messages should be written. It also affects the behavior of the <code>catopen()</code> function in determining the message catalog. Additional semantics of this variable, if any, are implementation-defined. The language and cultural conventions of diagnostic and informative messages whose format is unspecified by POSIX.1-202x should be affected by the setting of <code>LC_MESSAGES</code> .	
5505			
5506			
5507			
5508			
5509			
5510			
5511	<code>LC_MONETARY</code>		
5512		This variable shall determine the locale category for monetary-related numeric formatting information. Additional semantics of this variable, if any, are implementation-defined.	
5513			
5514			
5515	<code>LC_NUMERIC</code>		
5516		This variable shall determine the locale category for numeric formatting (for example, thousands separator and radix character) information in various utilities as well as the formatted I/O operations in <code>printf()</code> and <code>scanf()</code> and the string conversion functions in <code>strtod()</code> . Additional semantics of this variable, if any, are implementation-defined.	
5517			
5518			
5519			
5520			
5521	<code>LC_TIME</code>	This variable shall determine the locale category for date and time formatting information. It affects the behavior of the time functions in <code>strptime()</code> . Additional semantics of this variable, if any, are implementation-defined.	
5522			
5523			
5524	XSI <code>NLSPATH</code>	This variable shall contain a sequence of templates to be used by <code>catopen()</code> when attempting to locate message catalogs, and by the <code>gettext</code> family of functions when locating messages objects. Each template consists of an optional prefix, one or	
5525			
5526			

5527 more conversion specifications, and an optional suffix.

5528 The conversion specification descriptions below refer to a “currently active text
5529 domain”. The currently active text domain is, in decreasing order of precedence:

- 5530 • The *domain* parameter of the *gettext* family of functions or the *gettext* and
5531 *ngettext* utilities
- 5532 • The text domain bound by the last call to *textdomain()* when using a *gettext*
5533 family function, or the *TEXTDOMAIN* environment variable when using the
5534 *gettext* and *ngettext* utilities
- 5535 • The default text domain

5536 Conversion specifications consist of a '%' symbol, followed by a single-letter
5537 keyword. The following conversion specifications are currently defined:

5538 %N The value of the *name* parameter passed to *catopen()* or the currently active
5539 text domain of the *gettext* family of functions and the *gettext* and *ngettext*
5540 utilities (see above).

5541 %L The locale name given by the value of the active category (see *LANGUAGE*
5542 above) in either the current locale or, in the case of functions with an *_l* suffix,
5543 the provided locale object.

5544 %l The *language* element of the locale name that would result from a %L
5545 conversion.

5546 %t The *territory* element of the locale name that would result from a %L
5547 conversion.

5548 %c The *codeset* element of the locale name that would result from a %L conversion.

5549 %% A single '%' character.

5550 An empty string shall be substituted if the specified value is not currently defined.
5551 The separators <underscore> ('_') and <period> ('.') shall not be included in
5552 the %t and %c conversion specifications.

5553 Templates defined in *NLSPATH* are separated by <colon> characters (':'). A
5554 leading, trailing, or two adjacent <colon> characters ("::") shall be equivalent to
5555 specifying %N.

5556 Since <colon> is a separator in this context, directory names that might be used in
5557 *NLSPATH* should not include a <colon> character.

5558 Example 1, for an application that uses *catopen()* but does not use the *gettext* family
5559 of functions:

5560 `NLSPATH="/system/nlslib/%N.cat"`

5561 indicates that *catopen()* should look for all message catalogs in the directory
5562 `/system/nlslib`, where the catalog name should be constructed from the *name*
5563 argument (replacing %N) passed to *catopen()*, with the suffix `.cat`.

5564 Example 2, for an application that uses the *gettext* family of functions but does not
5565 use *catopen()*:

5566 `NLSPATH="/usr/lib/locale/fr/LC_MESSAGES/%N.mo"`

5567 indicates that the *gettext* family of functions (and the *gettext* and *ngettext* utilities)
5568 should look for all messages objects in the directory
5569 `/usr/lib/locale/fr/LC_MESSAGES`, where the messages object's name should be

5570 constructed from the currently active text domain (replacing %N), with the suffix
5571 **.mo**.

5572 Example 3, for an application that uses *catopen()* but does not use the *gettext* family
5573 of functions:

5574 `NLSPATH="%N.cat:/nlslib/%L/%N.cat"`

5575 indicates that *catopen()* should look for the requested message catalog in *name*,
5576 *name.cat*, and `/nlslib/locale/name.cat`, where *localename* is the locale name given
5577 by the value of the *LC_MESSAGES* category of the current locale.

5578 Example 4, for an application that uses the *gettext* family of functions but does not
5579 use *catopen()*:

5580 `NLSPATH="/usr/lib/locale/%L/%N.mo:/usr/lib/locale/fr/%N.mo"`

5581 indicates that the *gettext* family of functions (and the *gettext* and *ngettext* utilities)
5582 should look for all messages objects first in
5583 `/usr/lib/locale/localename/textdomain.mo`, and if not found there, then try in
5584 `/usr/lib/locale/fr/textdomain.mo`, where *localename* is the locale name given by the
5585 value of the active category in the current locale or provided locale object.

5586 Example 5, for an application that uses *catopen()* and the *gettext* family of
5587 functions:

5588 `NLSPATH="/usr/lib/locale/%L/%N.mo:/system/nlslib/%L/%N.cat"`

5589 indicates that the *gettext* family of functions (and the *gettext* and *ngettext* utilities)
5590 should look for all messages objects in `/usr/lib/locale/localename/textdomain.mo`,
5591 where *localename* is the locale name given by the value of the active category in the
5592 current locale or provided locale object. Also, *catopen()* should look for all message
5593 catalogs in the directory `/system/nlslib/locale/name.cat`, (assuming that
5594 `/usr/lib/locale/localename/name.mo` is not a message catalog). In this scenario,
5595 *catopen()* ignores all files that are not valid message catalogs while traversing
5596 *NLSPATH*. Furthermore, the *gettext* family of functions and the *gettext* and *ngettext*
5597 utilities ignore all files that are not valid messages objects found while traversing
5598 *NLSPATH*.

5599 Users should not set the *NLSPATH* variable unless they have a specific reason to
5600 override the default system path. Setting *NLSPATH* to override the default system
5601 path may produce undefined results in the standard utilities other than *gettext* and
5602 *ngettext*, and in applications with appropriate privileges.

5603 Specifying a relative pathname in the *NLSPATH* environment variable should be
5604 avoided without a specific reason, including the use of a leading, trailing, or two
5605 adjacent <colon> characters, since it may result in messages objects being searched
5606 for in a directory relative to the current working directory of the calling process; if
5607 the process calls the *chdir()* function, the directory searched for may also be
5608 changed.

5609 *TEXTDOMAIN*

5610 Specify the text domain name that the *gettext* and *ngettext* utilities use during the
5611 search for messages objects. This is identical to the messages object filename
5612 without the **.mo** suffix.

5613 *TEXTDOMAINDIR*

5614 Specify the pathname to the root directory of the messages object hierarchy the
5615 *gettext* and *ngettext* utilities use during the search for messages objects. If present, it

5616 XSI shall replace the default root directory pathname. *NLSPATH* has precedence over
5617 *TEXTDOMAINDIR*.

5618 The environment variables *LANG*, *LC_ALL*, *LC_COLLATE*, *LC_CTYPE*, *LC_MESSAGES*,
5619 *LC_MONETARY*, *LC_NUMERIC*, *LC_TIME*, and *NLSPATH* provide for the support of
5620 internationalized applications. The standard utilities shall make use of these environment
5621 variables as described in this section and the individual ENVIRONMENT VARIABLES sections
5622 for the utilities. If these variables specify locale categories that are not based upon the same
5623 underlying codeset, the results are unspecified.

5624 The values of locale categories shall be determined by a precedence order; the first condition met
5625 below determines the value:

- 5626 1. If the *LC_ALL* environment variable is defined and is not null, the value of *LC_ALL* shall
5627 be used.
- 5628 2. If the *LC_** environment variable (*LC_COLLATE*, *LC_CTYPE*, *LC_MESSAGES*,
5629 *LC_MONETARY*, *LC_NUMERIC*, *LC_TIME*) is defined and is not null, the value of the
5630 environment variable shall be used to initialize the category that corresponds to the
5631 environment variable.
- 5632 3. If the *LANG* environment variable is defined and is not null, the value of the *LANG*
5633 environment variable shall be used.
- 5634 4. If the *LANG* environment variable is not set or is set to the empty string, the
5635 implementation-defined default locale shall be used.

5636 If the locale value is "C" or "POSIX", the POSIX locale shall be used and the standard utilities
5637 behave in accordance with the rules in [Section 7.2](#) (on page 116) for the associated category.

5638 If the locale value begins with a <slash>, it shall be interpreted as the pathname of a file that was
5639 created in the output format used by the *localedef* utility; see OUTPUT FILES under *localedef*.
5640 Referencing such a pathname shall result in that locale being used for the indicated category.

5641 XSI If the locale value has the form:

```
5642 language[_territory] [.codeset]
```

5643 it refers to an implementation-provided locale, where settings of language, territory, and codeset
5644 are implementation-defined.

5645 *LC_COLLATE*, *LC_CTYPE*, *LC_MESSAGES*, *LC_MONETARY*, *LC_NUMERIC*, and *LC_TIME* are
5646 defined to accept an additional field *@modifier*, which allows the user to select a specific instance
5647 of localization data within a single category (for example, for selecting the dictionary as opposed
5648 to the character ordering of data). The syntax for these environment variables is thus defined as:

```
5649 [language[_territory] [.codeset] [@modifier]]
```

5650 For example, if a user wanted to interact with the system in French, but required to sort German
5651 text files, *LANG* and *LC_COLLATE* could be defined as:

```
5652 LANG=Fr_FR
```

```
5653 LC_COLLATE=De_DE
```

5654 This could be extended to select dictionary collation (say) by use of the *@modifier* field; for
5655 example:

```
5656 LC_COLLATE=De_DE@dict
```

5657 An implementation may support other formats.

9068	NAME		
9069		libintl.h — international messaging	+
9070	SYNOPSIS		+
9071		#include <libintl.h>	+
9072	DESCRIPTION		+
9073		The <libintl.h> header may define the macro TEXTDOMAINMAX. If defined, it shall have the	+
9074		same value as {TEXTDOMAIN_MAX} in <limits.h> .	+
9075		The following shall be declared as functions and may also be defined as macros. Function	+
9076		prototypes shall be provided.	+
9077		char *bindtextdomain(const char *, const char *);	+
9078		char *bind_textdomain_codeset(const char *, const char *);	+
9079		char *dcgettext(const char *, const char *, int);	+
9080		char *dcgettext_l(const char *, const char *, int, locale_t);	+
9081		char *dncgettext(const char *, const char *, const char *,	+
9082		unsigned long int, int);	+
9083		char *dncgettext_l(const char *, const char *, const char *,	+
9084		unsigned long int, int, locale_t);	+
9085		char *dgettext(const char *, const char *);	+
9086		char *dgettext_l(const char *, const char *, locale_t);	+
9087		char *dngettext(const char *, const char *, const char *,	+
9088		unsigned long int);	+
9089		char *dngettext_l(const char *, const char *, const char *,	+
9090		unsigned long int, locale_t);	+
9091		char *gettext(const char *);	+
9092		char *gettext_l(const char *, locale_t);	+
9093		char *ngettext(const char *, const char *, unsigned long int);	+
9094		char *ngettext_l(const char *, const char *,	+
9095		unsigned long int, locale_t);	+
9096		char *textdomain(const char *);	+
9097	APPLICATION USAGE		+
9098		None.	+
9099	RATIONALE		+
9100		Some historical implementations defined TEXTDOMAINMAX in this header. This standard	+
9101		instead defines {TEXTDOMAIN_MAX} in <limits.h> . This was done to allow the maximum	+
9102		length of a text domain name to vary depending on the filesystem type used to store message	+
9103		catalogs. Implementations are allowed to continue to define TEXTDOMAINMAX in this header	+
9104		as an extension to the standard (see XSH Section 2.2.2 , on page 467).	+
9105	FUTURE DIRECTIONS		+
9106		None.	+
9107	SEE ALSO		+
9108		XSH gettext , bindtextdomain()	+
9109	CHANGE HISTORY		+
9110		First released in Issue 8.	+
9111			+

9288 {PIPE_BUF}
 9289 Maximum number of bytes that is guaranteed to be atomic when writing to a pipe.
 9290 Minimum Acceptable Value: {_POSIX_PIPE_BUF}

9291 ADV {POSIX_ALLOC_SIZE_MIN}
 9292 Minimum number of bytes of storage actually allocated for any portion of a file.
 9293 Minimum Acceptable Value: Not specified.

9294 ADV {POSIX_REC_INCR_XFER_SIZE}
 9295 Recommended increment for file transfer sizes between the
 9296 {POSIX_REC_MIN_XFER_SIZE} and {POSIX_REC_MAX_XFER_SIZE} values.
 9297 Minimum Acceptable Value: Not specified.

9298 ADV {POSIX_REC_MAX_XFER_SIZE}
 9299 Maximum recommended file transfer size.
 9300 Minimum Acceptable Value: Not specified.

9301 ADV {POSIX_REC_MIN_XFER_SIZE}
 9302 Minimum recommended file transfer size.
 9303 Minimum Acceptable Value: Not specified.

9304 ADV {POSIX_REC_XFER_ALIGN}
 9305 Recommended file transfer buffer alignment.
 9306 Minimum Acceptable Value: Not specified.

9307 {SYMLINK_MAX}
 9308 Maximum number of bytes in a symbolic link.
 9309 Minimum Acceptable Value: {_POSIX_SYMLINK_MAX}

9310 {TEXTDOMAIN_MAX} +
 9311 Maximum length of a text domain name, not including the terminating null byte. +
 9312 Minimum Acceptable Value: {_POSIX_NAME_MAX} - 3 +
 9313 XSI {XOPEN_NAME_MAX} - 3

Runtime Inceasable Values

9314
 9315 The magnitude limitations in the following list shall be fixed by specific implementations. An
 9316 application should assume that the value of the symbolic constant defined by <limits.h> in a
 9317 specific implementation is the minimum that pertains whenever the application is run under
 9318 that implementation. A specific instance of a specific implementation may increase the value
 9319 relative to that supplied by <limits.h> for that implementation. The actual value supported by a
 9320 specific instance shall be provided by the *sysconf()* function.

9321 {BC_BASE_MAX}
 9322 Maximum *obase* values allowed by the *bc* utility.
 9323 Minimum Acceptable Value: {_POSIX2_BC_BASE_MAX}

9324 {BC_DIM_MAX}
 9325 Maximum number of elements permitted in an array by the *bc* utility.
 9326 Minimum Acceptable Value: {_POSIX2_BC_DIM_MAX}

9327 {BC_SCALE_MAX}
 9328 Maximum *scale* value allowed by the *bc* utility.
 9329 Minimum Acceptable Value: {_POSIX2_BC_SCALE_MAX}

9330 {BC_STRING_MAX}
 9331 Maximum length of a string constant accepted by the *bc* utility.
 9332 Minimum Acceptable Value: {_POSIX2_BC_STRING_MAX}

```

15290     _PC_PIPE_BUF
15291     _PC_PRIO_IO
15292     _PC_REC_INCR_XFER_SIZE
15293     _PC_REC_MAX_XFER_SIZE
15294     _PC_REC_MIN_XFER_SIZE
15295     _PC_REC_XFER_ALIGN
15296     _PC_SYMLINK_MAX
15297     _PC_SYNC_IO
15298     _PC_TEXTDOMAIN_MAX
15299     _PC_TIMESTAMP_RESOLUTION
15300     _PC_VDISABLE

```

+

15301 The **<unistd.h>** header shall define the following symbolic constants for *sysconf()*:

```

15302     _SC_2_C_BIND
15303     _SC_2_C_DEV
15304     _SC_2_CHAR_TERM
15305     _SC_2_FORT_RUN
15306     _SC_2_LOCALEDEF
15307     _SC_2_SW_DEV
15308     _SC_2_UPE
15309     _SC_2_VERSION
15310     _SC_ADVISORY_INFO
15311     _SC_AIO_LISTIO_MAX
15312     _SC_AIO_MAX
15313     _SC_AIO_PRIO_DELTA_MAX
15314     _SC_ARG_MAX
15315     _SC_ASYNCIO
15316     _SC_ATEXIT_MAX
15317     _SC_BARRIERS
15318     _SC_BC_BASE_MAX
15319     _SC_BC_DIM_MAX
15320     _SC_BC_SCALE_MAX
15321     _SC_BC_STRING_MAX
15322     _SC_CHILD_MAX
15323     _SC_CLK_TCK
15324     _SC_CLOCK_SELECTION
15325     _SC_COLL_WEIGHTS_MAX
15326     _SC_CPU_TIME
15327     _SC_DELAYTIMER_MAX
15328     _SC_EXPR_NEST_MAX
15329     _SC_FSYNC
15330     _SC_GETGR_R_SIZE_MAX
15331     _SC_GETPW_R_SIZE_MAX
15332     _SC_HOST_NAME_MAX
15333     _SC_IOV_MAX
15334     _SC_IPV6
15335     _SC_JOB_CONTROL
15336     _SC_LINE_MAX
15337     _SC_LOGIN_NAME_MAX
15338     _SC_MAPPED_FILES
15339     _SC_MEMLOCK
15340     _SC_MEMLOCK_RANGE

```

```

15484 pid_t      getpid(void);
15485 pid_t      getppid(void);
15486 XSI      int      getresgid(gid_t *, gid_t *, gid_t *);      +
15487 int      getresuid(uid_t *, uid_t *, uid_t *);      +
15488 pid_t      getsid(pid_t);
15489 uid_t      getuid(void);
15490 int      isatty(int);
15491 int      lchown(const char *, uid_t, gid_t);
15492 int      link(const char *, const char *);
15493 int      linkat(int, const char *, int, const char *, int);
15494 XSI      int      lockf(int, int, off_t);
15495 off_t     lseek(int, off_t, int);
15496 XSI      int      nice(int);
15497 long     pathconf(const char *, int);
15498 int      pause(void);
15499 int      pipe(int [2]);
15500 int      pipe2(int [2], int);
15501 int      posix_close(int, int);
15502 ssize_t  pread(int, void *, size_t, off_t);
15503 ssize_t  pwrite(int, const void *, size_t, off_t);
15504 ssize_t  read(int, void *, size_t);
15505 ssize_t  readlink(const char *restrict, char *restrict, size_t);
15506 ssize_t  readlinkat(int, const char *restrict, char *restrict, size_t);
15507 int      rmdir(const char *);
15508 int      setegid(gid_t);
15509 int      seteuid(uid_t);
15510 int      setgid(gid_t);
15511 int      setpgid(pid_t, pid_t);
15512 XSI      int      setregid(gid_t, gid_t);
15513 int      setresgid(gid_t, gid_t, gid_t);      +
15514 int      setresuid(uid_t, uid_t, uid_t);      +
15515 int      setreuid(uid_t, uid_t);
15516 pid_t    setsid(void);
15517 int      setuid(uid_t);
15518 unsigned sleep(unsigned);
15519 XSI      void     swab(const void *restrict, void *restrict, ssize_t);
15520 int      symlink(const char *, const char *);
15521 int      symlinkat(const char *, int, const char *);
15522 XSI      void     sync(void);
15523 long     sysconf(int);
15524 pid_t    tcgetpgrp(int);
15525 int      tcsetpgrp(int, pid_t);
15526 int      truncate(const char *, off_t);
15527 char     *ttyname(int);
15528 int      ttyname_r(int, char *, size_t);
15529 int      unlink(const char *);
15530 int      unlinkat(int, const char *, int);
15531 ssize_t  write(int, const void *, size_t);

15532 The <unistd.h> header shall declare the following external variables:

15533 extern char *optarg;
15534 extern int  opterr, optind, optopt;

```

	Header	Prefix	Suffix	Complete Name
16366	< aio.h>	aio_, lio_, AIO_, LIO_		
16367	<arpa/inet.h>	inet_		
16368	<ctype.h>	to[a-z], is[a-z]		
16369	<dlfcn.h>	RTLD_, dli_		
16370	<dirent.h>	d_, DT_		
16371	<fcntl.h>	l_		
16372	<fmtmsg.h>	MM_		
16373 XSI	<fnmatch.h>	FNM_		
16374	<ftw.h>	FTW		
16375 XSI	<glob.h>	gl_, GLOB_		
16376	<grp.h>	gr_		
16377	<libintl.h>			TEXTDOMAINMAX
16378	<limits.h>		_MAX, _MIN	
16379	<math.h>	M_		
16380 XSI	<mqueue.h>	mq_, MQ_		
16381 MSG	<ndbm.h>	dbm_, DBM_		
16382 XSI	<netdb.h>	ai_, h_, n_, p_, s_		
16383	<net/if.h>	if_, IF_		
16384	<netinet/in.h>	in_, ip_, s_, sin_, INADDR_, IPPROTO_		
16385				
16386 IP6	<netinet/tcp.h>	TCP_		
16387	<nl_types.h>	NL_		
16388	<poll.h>	pd_, ph_, ps_, POLL		
16389	<pthread.h>	pthread_, PTHREAD_		
16390	<pwd.h>	pw_		
16391	<regex.h>	re_, rm_, REG_		
16392	<sched.h>	sched_, SCHED_		
16393	<semaphore.h>	sem_, SEM_		
16394 CX	<signal.h>	sa_, si_, sigev_, sival_, uc_, BUS_, CLD_, FPE_, ILL_, SA_, SEGV_, SI_, SIGEV_		
16395				
16396 XSI		ss_, sv_, SS_, TRAP_		
16397	<stdlib.h>	str[a-z]		
16398	<string.h>	str[a-z], mem[a-z], wcs[a-z]		
16399 XSI	<sys/ipc.h>	ipc_, IPC_		key, pad, seq
16400	<sys/mman.h>	shm_, MAP_, MCL_, MS_, PROT_		
16401	<sys/msg.h>	msg, MSG_[A-Z]		msg
16402 XSI	<sys/resource.h>	rlim_, ru_, PRIO_, RLIMIT_, RUSAGE_		
16403				
16404	<sys/select.h>	fd_, fds_, FD_		
16405				

18104 2.9.5.2 Cancellation Points

18105 Cancellation points shall occur when a thread is executing the following functions:

18106	<i>accept()</i>	<i>nanosleep()</i>	<i>recvmsg()</i>
18107	<i>accept4()</i>	<i>open()</i>	<i>select()</i>
18108	<i>aio_suspend()</i>	<i>openat()</i>	<i>send()</i>
18109	<i>clock_nanosleep()</i>	<i>pause()</i>	<i>sendmsg()</i>
18110	<i>close()</i>	<i>poll()</i>	<i>sendto()</i>
18111	<i>connect()</i>	<i>ppoll()</i>	<i>sigsuspend()</i>
18112	<i>creat()</i>	<i>pread()</i>	<i>sigtimedwait()</i>
18113	<i>fcntl()†</i>	<i>pselect()</i>	<i>sigwait()</i>
18114	<i>fdatasync()</i>	<i>pthread_cond_clockwait()</i>	<i>sigwaitinfo()</i>
18115	<i>fsync()</i>	<i>pthread_cond_timedwait()</i>	<i>sleep()</i>
18116	<i>lockf()††</i>	<i>pthread_cond_wait()</i>	<i>tcdrain()</i>
18117	<i>mq_receive()</i>	<i>pthread_join()</i>	<i>wait()</i>
18118	<i>mq_send()</i>	<i>pthread_testcancel()</i>	<i>waitid()</i>
18119	<i>mq_timedreceive()</i>	<i>pwrite()</i>	<i>waitpid()</i>
18120	<i>mq_timedsend()</i>	<i>read()</i>	<i>write()</i>
18121	<i>msgrcv()</i>	<i>readv()</i>	<i>writew()</i>
18122	<i>msgsnd()</i>	<i>recv()</i>	
18123	<i>msync()</i>	<i>recvfrom()</i>	

18124 A cancellation point may also occur when a thread is executing the following functions:

18125	<i>access()</i>	<i>dngettext_l()</i>	<i>fgetws()</i>
18126	<i>bindtextdomain()</i>	<i>dprintf()</i>	<i>fntmsg()</i>
18127	<i>catclose()</i>	<i>endhostent()</i>	<i>fopen()</i>
18128	<i>catopen()</i>	<i>endnetent()</i>	<i>fpathconf()</i>
18129	<i>chmod()</i>	<i>endprotoent()</i>	<i>fprintf()</i>
18130	<i>chown()</i>	<i>endservent()</i>	<i>fputc()</i>
18131	<i>closedir()</i>	<i>faccessat()</i>	<i>fputs()</i>
18132	<i>closelog()</i>	<i>fchmod()</i>	<i>fputwc()</i>
18133	<i>ctermid()</i>	<i>fchmodat()</i>	<i>fputws()</i>
18134	<i>dcgettext()</i>	<i>fchown()</i>	<i>fread()</i>
18135	<i>dcgettext_l()</i>	<i>fchownat()</i>	<i>freopen()</i>
18136	<i>dcngettext()</i>	<i>fclose()</i>	<i>fscanf()</i>
18137	<i>dcngettext_l()</i>	<i>fcntl()†††</i>	<i>fseek()</i>
18138	<i>dgettext()</i>	<i>fflush()</i>	<i>fseeko()</i>
18139	<i>dgettext_l()</i>	<i>fgetc()</i>	<i>fsetpos()</i>
18140	<i>dlclose()</i>	<i>fgetpos()</i>	<i>fstat()</i>
18141	<i>dlopen()</i>	<i>fgets()</i>	<i>fstatat()</i>
18142	<i>dngettext()</i>	<i>fgetwc()</i>	<i>ftell()</i>

18143 † When the *cmd* argument is F_SETLKW.

18144 †† When the *function* argument is F_LOCK.

18145 ††† For any value of the *cmd* argument.

+

18146	<i>ftello()</i>	+	<i>mkstemp()</i>	<i>rewind()</i>
18147	<i>futimens()</i>		<i>mktime()</i>	<i>rewinddir()</i>
18148	<i>fwprintf()</i>		<i>ngettext()</i>	<i>scandir()</i>
18149	<i>fwrite()</i>		<i>ngettext_l()</i>	<i>scanf()</i>
18150	<i>fwscanf()</i>		<i>opendir()</i>	<i>seekdir()</i>
18151	<i>getaddrinfo()</i>		<i>openlog()</i>	<i>sem_clockwait()</i>
18152	<i>getc()</i>		<i>pathconf()</i>	<i>sem_timedwait()</i>
18153	<i>getc_unlocked()</i>		<i>perror()</i>	<i>sem_wait()</i>
18154	<i>getchar()</i>		<i>popen()</i>	<i>semop()</i>
18155	<i>getchar_unlocked()</i>		<i>posix_fadvise()</i>	<i>sethostent()</i>
18156	<i>getcwd()</i>		<i>posix_fallocate()</i>	<i>setnetent()</i>
18157	<i>getdelim()</i>		<i>posix_getdents()</i>	<i>setprotoent()</i>
18158	<i>getgrgid_r()</i>		<i>posix_madvise()</i>	<i>setservent()</i>
18159	<i>getgrnam_r()</i>		<i>posix_openpt()</i>	<i>stat()</i>
18160	<i>gethostid()</i>		<i>posix_spawn()</i>	<i>strerror_l()</i>
18161	<i>gethostname()</i>		<i>posix_spawnnp()</i>	<i>strerror_r()</i>
18162	<i>getline()</i>		<i>posix_typed_mem_open()</i>	<i>strftime()</i>
18163	<i>getlogin_r()</i>		<i>printf()</i>	<i>strftime_l()</i>
18164	<i>getnameinfo()</i>		<i>psiginfo()</i>	<i>symlink()</i>
18165	<i>getpwnam_r()</i>		<i>psignal()</i>	<i>symlinkat()</i>
18166	<i>getpwuid_r()</i>		<i>pthread_rwlock_clockrdlock()</i>	<i>sync()</i>
18167	<i>gettext()</i>		<i>pthread_rwlock_clockwrlock()</i>	<i>syslog()</i>
18168	<i>gettext_l()</i>		<i>pthread_rwlock_rdlock()</i>	<i>tmpfile()</i>
18169	<i>getwc()</i>		<i>pthread_rwlock_timedrdlock()</i>	<i>tmpnam()</i>
18170	<i>getwchar()</i>		<i>pthread_rwlock_timedwrlock()</i>	<i>ttyname_r()</i>
18171	<i>glob()</i>		<i>pthread_rwlock_wrlock()</i>	<i>tzset()</i>
18172	<i>iconv_close()</i>		<i>ptsname()</i>	<i>ungetc()</i>
18173	<i>iconv_open()</i>		<i>ptsname_r()</i>	<i>ungetwc()</i>
18174	<i>link()</i>		<i>putc()</i>	<i>unlink()</i>
18175	<i>linkat()</i>		<i>putc_unlocked()</i>	<i>unlinkat()</i>
18176	<i>lio_listio()</i>		<i>putchar()</i>	<i>utimensat()</i>
18177	<i>localtime_r()</i>		<i>putchar_unlocked()</i>	<i>utimes()</i>
18178	<i>lockf()</i>		<i>puts()</i>	<i>vdprintf()</i>
18179	<i>lseek()</i>		<i>putwc()</i>	<i>vfprintf()</i>
18180	<i>lstat()</i>		<i>putwchar()</i>	<i>vfwprintf()</i>
18181	<i>mkdir()</i>		<i>readdir_r()</i>	<i>vprintf()</i>
18182	<i>mkdirat()</i>		<i>readlink()</i>	<i>vfprintf()</i>
18183	<i>mkdtemp()</i>		<i>readlinkat()</i>	<i>wcsftime()</i>
18184	<i>mkfifo()</i>		<i>remove()</i>	<i>wordexp()</i>
18185	<i>mkfifoat()</i>		<i>rename()</i>	<i>wprintf()</i>
18186	<i>mknod()</i>		<i>renameat()</i>	<i>wscanf()</i>
18187	<i>mknodat()</i>			

18188 In addition, a cancellation point may occur when a thread is executing any function that this
 18189 standard does not require to be thread-safe but the implementation documents as being thread-
 18190 safe. If a thread is cancelled while executing a non-thread-safe function, the behavior is
 18191 undefined.

18192 An implementation shall not introduce cancellation points into any other functions specified in
 18193 this volume of POSIX.1-202x.

18194 The side-effects of acting upon a cancellation request while suspended during a call of a function
 18195 are the same as the side-effects that may be seen in a single-threaded program when a call to a
 18196 function is interrupted by a signal and the given function returns [EINTR]. Any such side-

21422	NAME		
21423		bindtextdomain, bind_textdomain_codeset, textdomain — text domain manipulation functions	+
21424	SYNOPSIS		+
21425		#include <libintl.h>	+
21426		char *bindtextdomain(const char *domainname, const char *dirname);	+
21427		char *bind_textdomain_codeset(const char *domainname,	+
21428		const char *codeset);	+
21429		char *textdomain(const char *domainname);	+
21430	DESCRIPTION		+
21431		The <i>textdomain()</i> function shall set or query the name of the current text domain of the calling	+
21432		process. The application shall ensure that the <i>domainname</i> argument is either a null pointer	+
21433		(when querying), an empty string, or a string that, when used by the <i>gettext</i> family of functions	+
21434		to construct a pathname to a messages object, results in a valid pathname. For portable	+
21435		applications, it should only contain characters from the portable filename character set.	+
21436		The text domain setting made by the last successful call to <i>textdomain()</i> shall remain in effect	+
21437		across subsequent calls to <i>setlocale()</i> , <i>uselocale()</i> , and the <i>gettext</i> family of functions.	+
21438		Applications should not use text domains whose names begin with the strings "SYS_" or	+
21439		"libc". These prefixes are reserved for implementation use.	+
21440		The current setting of the text domain can be queried without affecting the current state of the	+
21441		domain by calling <i>textdomain()</i> with <i>domainname</i> set to a null pointer. Calling <i>textdomain()</i> with a	+
21442		<i>domainname</i> argument of an empty string shall set the text domain to the default domain,	+
21443		"messages".	+
21444		The <i>bindtextdomain()</i> function shall set or query the binding of a text domain to a <i>dirname</i> that is	+
21445		used by the <i>gettext</i> family of functions to construct a pathname to a messages object in the text	+
21446		domain:	+
21447		• If <i>domainname</i> is a null pointer or an empty string, <i>bindtextdomain()</i> shall make no changes	+
21448		and return a null pointer without changing <i>errno</i> .	+
21449		• Otherwise, if <i>dirname</i> is a non-empty string:	+
21450		— If <i>domainname</i> is not already bound, <i>bindtextdomain()</i> shall bind the text domain	+
21451		specified by <i>domainname</i> to the pathname pointed to by <i>dirname</i> and return the bound	+
21452		directory pathname on success or a null pointer on failure.	+
21453		— If <i>domainname</i> is already bound, <i>bindtextdomain()</i> shall replace the existing binding	+
21454		with the pathname pointed to by <i>dirname</i> and return the bound directory pathname	+
21455		on success or a null pointer on failure. On failure, the existing binding shall remain	+
21456		unchanged.	+
21457		It is unspecified whether the <i>bindtextdomain()</i> function performs pathname resolution on	+
21458		<i>dirname</i> , or whether that is done by the <i>gettext</i> family of functions.	+
21459		• Otherwise, if <i>dirname</i> is a null pointer:	+
21460		— If <i>domainname</i> is bound, the function shall return the bound directory pathname.	+
21461		— If <i>domainname</i> is not bound, the function shall return the implementation-defined	+
21462		default directory pathname used by the <i>gettext</i> family of functions.	+
21463		• Otherwise, <i>dirname</i> is an empty string and the behavior is unspecified.	+
21464		If a text domain is bound to a relative pathname and the current working directory is changed	+
21465		after the binding is established, the pathnames used by the <i>gettext</i> family of functions to locate	+

21466	messages objects for that text domain are unspecified.	+
21467	The <i>bind_textdomain_codeset()</i> function shall set or query the binding of a text domain to the	+
21468	output codeset used by the <i>gettext</i> family of functions for message strings retrieved from	+
21469	messages objects for the text domain specified by <i>domainname</i> :	+
21470	• If <i>domainname</i> is a null pointer or an empty string, <i>bind_textdomain_codeset()</i> shall make no	+
21471	changes and return a null pointer without changing <i>errno</i> .	+
21472	• Otherwise, if <i>codeset</i> is a non-empty string:	+
21473	— If <i>domainname</i> is not already bound, <i>bind_textdomain_codeset()</i> shall bind the text	+
21474	domain specified by <i>domainname</i> to the codeset pointed to by <i>codeset</i> and return the	+
21475	newly bound codeset on success or a null pointer on failure.	+
21476	— If <i>domainname</i> is already bound, <i>bind_textdomain_codeset()</i> shall replace the existing	+
21477	binding with the codeset pointed to by <i>codeset</i> and return the newly bound codeset	+
21478	on success or a null pointer on failure. On failure, the existing binding shall remain	+
21479	unchanged.	+
21480	The application shall ensure that the <i>codeset</i> argument, if non-empty, is a valid codeset	+
21481	name that can be used as the <i>tocode</i> argument of the <i>iconv_open()</i> function, and that in the	+
21482	codeset it specifies, the <NUL> character corresponds to a single null byte.	+
21483	• Otherwise, if <i>codeset</i> is a null pointer:	+
21484	— If <i>domainname</i> is bound, the function shall return the bound codeset.	+
21485	— If <i>domainname</i> is not bound, the function shall return the implementation-defined	+
21486	default codeset used by the <i>gettext</i> family of functions.	+
21487	• Otherwise, <i>codeset</i> is an empty string and the behavior is unspecified.	+
21488	If <i>codeset</i> is a null pointer and <i>domainname</i> is a non-empty string, <i>bind_textdomain_codeset()</i> shall	+
21489	return the current codeset for the named domain, or a null pointer if a codeset has not yet been	+
21490	set. The <i>bind_textdomain_codeset()</i> function can be called multiple times. If successfully called	+
21491	multiple times with the same <i>domainname</i> argument, the last such call shall override the setting	+
21492	made by the previous such call.	+
21493	RETURN VALUE	+
21494	The return value from a successful <i>textdomain()</i> call shall be a pointer to a string containing the	+
21495	current setting of the text domain. If <i>domainname</i> is a null pointer, <i>textdomain()</i> shall return a	+
21496	pointer to the string containing the current text domain. If <i>textdomain()</i> was not previously	+
21497	called and <i>domainname</i> is a null string, the name of the default text domain shall be returned.	+
21498	The name of the default text domain shall be the string "messages". If <i>textdomain()</i> fails, a null	+
21499	pointer shall be returned and <i>errno</i> shall be set to indicate the error.	+
21500	For <i>bindtextdomain()</i> return values see the DESCRIPTION. When <i>bindtextdomain()</i> is called with	+
21501	a non-empty <i>domainname</i> and returns a null pointer, it shall set <i>errno</i> to indicate the error. When	+
21502	<i>bindtextdomain()</i> returns a pathname for a bound text domain, the return value shall be a pointer	+
21503	to a copy of the <i>dirname</i> string passed to the <i>bindtextdomain()</i> call that created the binding. The	+
21504	returned string shall remain valid until the next successful call to <i>bindtextdomain()</i> with a non-	+
21505	empty <i>dirname</i> and same <i>domainname</i> . The application shall ensure that it does not modify the	+
21506	returned string.	+
21507	A call to the <i>bind_textdomain_codeset()</i> function with a non-empty <i>domainname</i> argument shall	+
21508	return one of the following:	+

21509	• The currently bound codeset name for that text domain if <i>codeset</i> is a null pointer	+
21510	• The newly bound codeset if <i>codeset</i> is non-empty	+
21511	• A null pointer without changing <i>errno</i> if no codeset has yet been bound for that text	+
21512	domain	+
21513	The application shall ensure that it does not modify the returned string. A subsequent call to	+
21514	<i>bind_textdomain_codeset()</i> with a non-empty <i>domainname</i> argument might invalidate the returned	+
21515	pointer or overwrite the string content. The returned pointer might also be invalidated if the	+
21516	calling thread is terminated. If <i>bind_textdomain_codeset()</i> fails, a null pointer shall be returned	+
21517	and <i>errno</i> shall be set to indicate the error.	+
21518	ERRORS	+
21519	For the conditions under which <i>bindtextdomain()</i> —if it performs pathname resolution—fails and	+
21520	may fail, refer to <i>open()</i> .	+
21521	In addition, the <i>textdomain()</i> , <i>bindtextdomain()</i> , and <i>bind_textdomain_codeset()</i> functions may fail	+
21522	if:	+
21523	[ENOMEM] Insufficient memory available.	+
21524	EXAMPLES	+
21525	See the examples for <i>gettext</i> .	+
21526	APPLICATION USAGE	+
21527	A text <i>domainname</i> is limited to {TEXTDOMAIN_MAX} bytes.	+
21528	Application developers are responsible for ensuring that the text domain used is not used by	+
21529	other applications. To minimize the chances of collision, developers can prefix text domains with	+
21530	their company or application name (or both) and an underscore. For example, if your	+
21531	application name was "foo" and you wanted to use the text domain "errors", you could	+
21532	instead use the text domain "foo_errors". Note that if an application can be installed with a	+
21533	configurable name, a text domain prefix based on the application name should change with the	+
21534	application name.	+
21535	Specifying a relative pathname to the <i>bindtextdomain()</i> function should be avoided, since it may	+
21536	result in messages objects being searched for in a directory relative to the current working	+
21537	directory of the calling process; if the process calls the <i>chdir()</i> function, the directory searched for	+
21538	may also be changed.	+
21539	Since pathname resolution of <i>dirname</i> might not be performed by <i>bindtextdomain()</i> , but could be	+
21540	performed later by the <i>gettext</i> family of functions, and since the latter have no way to report an	+
21541	error, applications should verify, using for example <i>stat()</i> , that the directory is accessible if this is	+
21542	desired.	+
21543	RATIONALE	+
21544	Although the return type of these functions ought to be const char * , it is char * to match	+
21545	historical practice.	+
21546	Pathname resolution of the <i>dirname</i> argument passed to <i>bindtextdomain()</i> may be performed by	+
21547	<i>bindtextdomain()</i> itself or by the <i>gettext</i> family of functions. If pathname resolution fails in one of	+
21548	the <i>gettext</i> family of functions, it is neither allowed to modify <i>errno</i> nor to return an error, but if	+
21549	pathname resolution fails in <i>bindtextdomain()</i> , it is required to report an error and set <i>errno</i> just	+
21550	like <i>open()</i> does.	+
21551	Historically, <i>bindtextdomain()</i> did not perform pathname resolution. However, the standard	+
21552	developers decided to allow this as an option so that future implementations can, if desired,	+
21553	open a file descriptor for that directory in <i>bindtextdomain()</i> and then use that file descriptor with	+

21554	<i>openat()</i> in the <i>gettext</i> family of functions.	+
21555	The <i>dirname</i> parameter to <i>bindtextdomain()</i> may need to be copied to avoid the possibility of the application releasing the memory used by the argument while the <i>gettext</i> family of functions may still need to reference it.	+
21556		+
21557		+
21558		+
21559	When <i>bindtextdomain()</i> is called with a non-empty <i>domainname</i> and an empty <i>dirname</i> , historical implementations of the <i>gettext</i> family of functions use the empty string for the <i>dirname</i> part of the messages object pathname, resulting in an absolute pathname of the form <i>/localename/categoryname/textdomainname.mo</i> . The standard developers did not believe this behavior to be useful. Using the empty <i>dirname</i> case as a way to remove an existing binding seemed to be a more useful behavior, and would be consistent with the behavior of <i>textdomain()</i> . However, because no historical implementations behave this way, the behavior is left unspecified.	+
21560		+
21561		+
21562		+
21563		+
21564		+
21565		+
21566		+
21567	Some implementations set <i>errno</i> to [EAGAIN] to signal memory allocation failures that might succeed if retried and [ENOMEM] for failures that are unlikely to ever succeed, for example due to configured limits. Section 2.3 (on page 475) permits this behavior; when multiple error conditions are simultaneously true there is no precedence between them.	+
21568		+
21569		+
21570		+
21570	FUTURE DIRECTIONS	+
21571	A future version of this standard may require implementations to prefix implementation-provided text domains with either "SYS_" or a prefix related to the implementor's company name to avoid namespace collisions.	+
21572		+
21573		+
21574	A future version of this standard may require <i>bindtextdomain()</i> to remove any binding for <i>domainname</i> when called with a non-empty <i>domainname</i> and an empty <i>dirname</i> .	+
21575		+
21576	SEE ALSO	+
21577	<i>gettext</i> , <i>iconv_open()</i> , <i>setlocale()</i> , <i>uselocale()</i>	+
21578	XBD <libintl.h> , <limits.h>	+
21579	XCU <i>msgfmt</i> , <i>xgettext</i>	+
21580	CHANGE HISTORY	+
21581	First released in Issue 8.	+
21582		+

22136 **NAME**

22137 catopen — open a message catalog

22138 **SYNOPSIS**

22139 #include <nl_types.h>

22140 nl_catd catopen(const char *name, int oflag);

22141 **DESCRIPTION**

22142 The *catopen()* function shall open a message catalog and return a message catalog descriptor.
 22143 The *name* argument specifies the name of the message catalog to be opened. If *name* contains a
 22144 `'/'`, then *name* specifies a pathname for the message catalog. Otherwise, the environment
 22145 variable *NLSPATH* is used with *name* substituted for the `%N` conversion specification (see XBD
 22146 [Chapter 8](#), on page 155); if *NLSPATH* exists in the environment when the process starts, then if
 22147 the process has appropriate privileges, the behavior of *catopen()* is undefined. If *NLSPATH* does
 22148 not exist in the environment, or if a message catalog cannot be found in any of the components
 22149 specified by *NLSPATH*, then an implementation-defined default path shall be used. This default
 22150 may be affected by the setting of *LC_MESSAGES* if the value of *oflag* is *NL_CAT_LOCALE*, or
 22151 XSI the *LANG* environment variable if *oflag* is 0. When searching *NLSPATH*, *catopen()* shall ignore +
 22152 any files it finds that are not valid message catalog files.

22153 A message catalog descriptor shall remain valid in a process until that process closes it, or a
 22154 successful call to one of the *exec* functions. A change in the setting of the *LC_MESSAGES*
 22155 category may invalidate existing open catalogs.

22156 If a file descriptor is used to implement message catalog descriptors, the *FD_CLOEXEC* flag
 22157 shall be set; see <*fcntl.h*>.

22158 If the value of the *oflag* argument is 0, the *LANG* environment variable is used to locate the
 22159 catalog without regard to the *LC_MESSAGES* category. If the *oflag* argument is
 22160 *NL_CAT_LOCALE*, the *LC_MESSAGES* category is used to locate the message catalog (see XBD
 22161 [Section 8.2](#), on page 157).

22162 **RETURN VALUE**

22163 Upon successful completion, *catopen()* shall return a message catalog descriptor for use on
 22164 subsequent calls to *catgets()* and *catclose()*. Otherwise, *catopen()* shall return (*nl_catd*) -1 and set
 22165 *errno* to indicate the error.

22166 **ERRORS**22167 The *catopen()* function may fail if:

22168 [EACCES] Search permission is denied for the component of the path prefix of the
 22169 message catalog or read permission is denied for the message catalog.

22170 [EMFILE] All file descriptors available to the process are currently open.

22171 [ENAMETOOLONG]

22172 The length of a component of a pathname is longer than {NAME_MAX}.

22173 [ENAMETOOLONG]

22174 The length of a pathname exceeds {PATH_MAX}, or pathname resolution of a
 22175 symbolic link produced an intermediate result with a length that exceeds
 22176 {PATH_MAX}.

22177 [ENFILE] Too many files are currently open in the system.

22178 [ENOENT] The *name* argument contains a `'/'` and does not name an existing message
 22179 XSI catalog, the *name* argument does not contain a `'/'` and searching *NLSPATH* (if
 22180 set) and then the implementation-defined default path for a message catalog

22181 with that name failed, one or more files exist but all are of an invalid format,
 22182 or the *name* argument points to an empty string.

22183 [ENOMEM] Insufficient storage space is available.

22184 [ENOTDIR] A component of the path prefix of the message catalog names an existing file
 22185 that is neither a directory nor a symbolic link to a directory, or the pathname
 22186 of the message catalog contains at least one non-`<slash>` character and ends
 22187 with one or more trailing `<slash>` characters and the last pathname
 22188 component names an existing file that is neither a directory nor a symbolic
 22189 link to a directory.

22190 EXAMPLES

22191 None.

22192 APPLICATION USAGE

22193 Some implementations of *catopen()* use *malloc()* to allocate space for internal buffer areas. The
 22194 *catopen()* function may fail if there is insufficient storage space available to accommodate these
 22195 buffers.

22196 Conforming applications must assume that message catalog descriptors are not valid after a call
 22197 to one of the *exec* functions.

22198 Application developers should be aware that guidelines for the location of message catalogs
 22199 have not yet been developed. Therefore they should take care to avoid conflicting with catalogs
 22200 used by other applications and the standard utilities.

22201 To be sure that messages produced by an application running with appropriate privileges cannot
 22202 be used by an attacker setting an unexpected value for *NLSPATH* in the environment to confuse
 22203 a system administrator, such applications should use pathnames containing a `'/'` to get defined
 22204 behavior when using *catopen()* to open a message catalog.

22205 RATIONALE

22206 None.

22207 FUTURE DIRECTIONS

22208 None.

22209 SEE ALSO

22210 [catclose\(\)](#), [catgets\(\)](#)

22211 XBD Chapter 8 (on page 155), [<fcntl.h>](#), [<nl_types.h>](#),

22212 CHANGE HISTORY

22213 First released in Issue 2.

22214 Issue 7

22215 Austin Group Interpretation 1003.1-2001 #143 is applied.

22216 SD5-XBD-ERN-4 is applied, changing the definition of the [EMFILE] error.

22217 The *catopen()* function is moved from the XSI option to the Base.

22218 POSIX.1-2008, Technical Corrigendum 1, XSH/TC1-2008/0045 [324] is applied.

22219 POSIX.1-2008, Technical Corrigendum 2, XSH/TC2-2008/0054 [645], XSH/TC2-2008/0055 [497],
 22220 and XSH/TC2-2008/0056 [497] are applied.

25023 **NAME**

25024 dcgettext, dcgettext_l, dcngettext, dcngettext_l, dgettext, dgettext_l — message handling
25025 functions

25026 **SYNOPSIS**

```
25027 #include <libintl.h>
25028 char *dcgettext(const char *domainname, const char *msgid,
25029               int category);
25030 char *dcgettext_l(const char *domainname, const char *msgid,
25031                 int category, locale_t locale);
25032 char *dcngettext(const char *domainname, const char *msgid,
25033                 const char *msgid_plural, unsigned long int n,
25034                 int category);
25035 char *dcngettext_l(const char *domainname, const char *msgid,
25036                  const char *msgid_plural, unsigned long int n,
25037                  int category, locale_t locale);
25038 char *dgettext(const char *domainname, const char *msgid);
25039 char *dgettext_l(const char *domainname, const char *msgid,
25040                 locale_t locale);
```

25041 **DESCRIPTION**

25042 Refer to [gettext](#).

25654 **NAME**

25655 dngettext, dngettext_l — message handling functions

25656 **SYNOPSIS**

25657 #include <libintl.h>

```
25658 char *dngettext(const char *domainname, const char *msgid,  
25659                const char *msgid_plural, unsigned long int n);  
25660 char *dngettext_l(const char *domainname, const char *msgid,  
25661                  const char *msgid_plural, unsigned long int n,  
25662                  locale_t locale);
```

25663 **DESCRIPTION**25664 Refer to [gettext](#).

30747 **NAME**

30748 `fpathconf`, `pathconf` — get configurable pathname variables

30749 **SYNOPSIS**

```
30750 #include <unistd.h>
30751 long fpathconf(int fildes, int name);
30752 long pathconf(const char *path, int name);
```

30753 **DESCRIPTION**

30754 The `fpathconf()` and `pathconf()` functions shall determine the current value of a configurable limit
30755 or option (*variable*) that is associated with a file or directory.

30756 For `pathconf()`, the *path* argument points to the pathname of a file or directory.

30757 For `fpathconf()`, the *filde*s argument is an open file descriptor.

30758 The *name* argument represents the variable to be queried relative to that file or directory.
30759 Implementations shall support all of the variables listed in the following table and may support
30760 others. The variables in the following table come from `<limits.h>` or `<unistd.h>` and the
30761 symbolic constants, defined in `<unistd.h>`, are the corresponding values used for *name*.

Variable	Value of <i>name</i>	Requirements
{FILESIZEBITS}	_PC_FILESIZEBITS	4,7
{LINK_MAX}	_PC_LINK_MAX	1
{MAX_CANON}	_PC_MAX_CANON	2
{MAX_INPUT}	_PC_MAX_INPUT	2
{NAME_MAX}	_PC_NAME_MAX	3,4
{PATH_MAX}	_PC_PATH_MAX	4,5
{PIPE_BUF}	_PC_PIPE_BUF	6
{POSIX2_SYMLINKS}	_PC_2_SYMLINKS	4
{POSIX_ALLOC_SIZE_MIN}	_PC_ALLOC_SIZE_MIN	10
{POSIX_REC_INCR_XFER_SIZE}	_PC_REC_INCR_XFER_SIZE	10
{POSIX_REC_MAX_XFER_SIZE}	_PC_REC_MAX_XFER_SIZE	10
{POSIX_REC_MIN_XFER_SIZE}	_PC_REC_MIN_XFER_SIZE	10
{POSIX_REC_XFER_ALIGN}	_PC_REC_XFER_ALIGN	10
{SYMLINK_MAX}	_PC_SYMLINK_MAX	4,9
{TEXTDOMAIN_MAX}	_PC_TEXTDOMAIN_MAX	3,4
_POSIX_CHOWN_RESTRICTED	_PC_CHOWN_RESTRICTED	7
_POSIX_NO_TRUNC	_PC_NO_TRUNC	3,4
_POSIX_VDISABLE	_PC_VDISABLE	2
_POSIX_ASYNC_IO	_PC_ASYNC_IO	8
_POSIX_FALLOC	_PC_FALLOC	8
_POSIX_PRIO_IO	_PC_PRIO_IO	8
_POSIX_SYNC_IO	_PC_SYNC_IO	8
_POSIX_TIMESTAMP_RESOLUTION	_PC_TIMESTAMP_RESOLUTION	1

37100 **NAME**

37101 getresgid — get real group ID, effective group ID, and saved set-group-ID

37102 **SYNOPSIS**

37103 XSI #include <unistd.h>

37104 int getresgid(gid_t *rgid, gid_t *egid, gid_t *sgid);

37105

37106 **DESCRIPTION**

37107 The *getresgid()* function shall store the real group ID, effective group ID, and saved set-group-ID
37108 of the calling process in the locations pointed to by the arguments *rgid*, *egid*, and *sgid*,
37109 respectively.

37110 **RETURN VALUE**

37111 Upon successful completion, 0 shall be returned. Otherwise, -1 shall be returned and *errno* set to
37112 indicate the error.

37113 **ERRORS**

37114 No errors are defined.

37115 **EXAMPLES**

37116 None.

37117 **APPLICATION USAGE**

37118 None.

37119 **RATIONALE**

37120 None.

37121 **FUTURE DIRECTIONS**

37122 None.

37123 **SEE ALSO**

37124 *exec*, *getegid()*, *geteuid()*, *getgid()*, *getresuid()*, *getuid()*, *setegid()*, *seteuid()*, *setgid()*, *setregid()*,
37125 *setresgid()*, *setresuid()*, *setreuid()*, *setuid()*

37126 XBD <unistd.h>

37127 **CHANGE HISTORY**

37128 First released in Issue 8.

37129

NAME

37130 getresuid — get real user ID, effective user ID, and saved set-user-ID
37131

SYNOPSIS

```
37133 XSI #include <unistd.h>  
37134 int getresuid(uid_t *ruid, uid_t *euid, uid_t *suid);  
37135
```

DESCRIPTION

37136 The *getresuid()* function shall store the real user ID, effective user ID, and saved set-user-ID of
37137 the calling process in the locations pointed to by the arguments *ruid*, *euid*, and *suid*, respectively.
37138

RETURN VALUE

37139 Upon successful completion, 0 shall be returned. Otherwise, -1 shall be returned and *errno* set to
37140 indicate the error.
37141

ERRORS

37142 No errors are defined.
37143

EXAMPLES

37144 None.
37145

APPLICATION USAGE

37146 None.
37147

RATIONALE

37148 None.
37149

FUTURE DIRECTIONS

37150 None.
37151

SEE ALSO

37152 *exec*, *getegid()*, *geteuid()*, *getgid()*, *getresgid()*, *getuid()*, *setegid()*, *seteuid()*, *setgid()*, *setregid()*,
37153 *setresgid()*, *setresuid()*, *setreuid()*, *setuid()*
37154

37155 XBD <unistd.h>

CHANGE HISTORY

37156 First released in Issue 8.
37157
37158

37628 **NAME**

37629 `dgettext`, `dgettext_l`, `dcgettext`, `dcgettext_l`, `gettext`, `gettext_l`, `ngettext`, `ngettext_l`, `dngettext`,
 37630 `dngettext_l`, `dcngettext`, `dcngettext_l` — message handling functions

37631 **SYNOPSIS**

```
37632 #include <libintl.h>
37633 char *dgettext(const char *domainname, const char *msgid);
37634 char *dgettext_l(const char *domainname, const char *msgid,
37635                 locale_t locale);
37636 char *dcgettext(const char *domainname, const char *msgid,
37637                int category);
37638 char *dcgettext_l(const char *domainname, const char *msgid,
37639                  int category, locale_t locale);
37640 char *dngettext(const char *domainname, const char *msgid,
37641                const char *msgid_plural, unsigned long int n);
37642 char *dngettext_l(const char *domainname, const char *msgid,
37643                  const char *msgid_plural, unsigned long int n,
37644                  locale_t locale);
37645 char *dcngettext(const char *domainname, const char *msgid,
37646                 const char *msgid_plural, unsigned long int n,
37647                 int category);
37648 char *dcngettext_l(const char *domainname, const char *msgid,
37649                   const char *msgid_plural, unsigned long int n,
37650                   int category, locale_t locale);
37651 char *gettext(const char *msgid);
37652 char *gettext_l(const char *msgid, locale_t locale);
37653 char *ngettext(const char *msgid, const char *msgid_plural,
37654               unsigned long int n);
37655 char *ngettext_l(const char *msgid, const char *msgid_plural,
37656                 unsigned long int n, locale_t locale);
```

37657 **DESCRIPTION**

37658 The `gettext()` function shall:

- 37659 • attempt to locate a suitable messages object (described in detail below) for the
 37660 `LC_MESSAGES` category in the current locale, and for the current text domain (see
 37661 `bindtextdomain()`), containing the string identified by `msgid`,
- 37662 • retrieve the string identified by `msgid` from the messages object,
- 37663 • convert the string to the output codeset if necessary (described in detail below), and
- 37664 • return the result.

37665 If the locale name in effect is "POSIX" or "C" (i.e. the name associated with the `LC_MESSAGES`
 37666 locale category in the current locale), or if no suitable messages object exists, or if no string
 37667 identified by `msgid` exists in the messages object, or if an error occurs, `msgid` shall be returned.

37668 The `dgettext()` function shall be equivalent to `gettext()`, except `domainname` shall be used instead
 37669 of the current text domain to locate the messages object.

37670 The `dcgettext()` function shall be equivalent to `dgettext()`, except the locale category identified by
 37671 `category` shall be used instead of `LC_MESSAGES`.

37672 The `ngettext()` function shall be equivalent to `gettext()`, except:

- 37673 • The string to retrieve shall be identified by a combination of *msgid* and *n* (see *msgfmt*).
- 37674 • If the locale name in effect is "POSIX" or "C", or if no suitable messages object exists, or if
- 37675 no string identified by the combination of *msgid* and *n* exists in the messages object, or if an
- 37676 error occurs, the return value shall be *msgid* if *n* is 1, otherwise *msgid_plural*.

37677 The *dngettext()* function shall be equivalent to *ngettext()*, except *domainname* shall be used

37678 instead of the current text domain to locate the messages object.

37679 The *dcngettext()* function shall be equivalent to *dngettext()*, except the locale category identified

37680 by *category* shall be used instead of *LC_MESSAGES*.

37681 The **_l()* functions shall be equivalent to their counterparts without the *_l* suffix, except *locale*

37682 shall be used instead of the current locale. If *locale* is the special locale object

37683 *LC_GLOBAL_LOCALE* or is not a valid locale object handle, the behavior is undefined.

37684 The application shall ensure that the *msgid* and *msgid_plural* arguments are strings. If either

37685 *msgid* or *msgid_plural* is an empty string, or contains characters not in the portable character set,

37686 the results are unspecified. If the *category* argument is *LC_ALL*, the results are unspecified.

37687 The location of the messages object shall be determined according to the following criteria,

37688 stopping when the first messages object is found:

- 37689 XSI 1. If the *NLSPATH* environment variable is set to a non-empty string, an *NLSPATH* search
- 37690 shall be performed as described in XBD Section 8.2 (on page 157). If *NLSPATH* identifies
- 37691 more than one template to use, each template in turn shall be used until a valid messages
- 37692 object is found.
- 37693 2. If the *LANGUAGE* environment variable is set to a non-empty string, a *LANGUAGE*
- 37694 search shall be performed as described below. If *LANGUAGE* identifies more than one
- 37695 directory to search, each directory shall be searched until a valid messages object is found.
- 37696 3. A single-locale search shall be performed as described below.

37697 XSI For the *NLSPATH* search and the single-locale search, the single locale name used to locate the

37698 messages object shall be the locale name associated with the selected locale category from the

37699 current locale, or the provided locale object if calling one of the **_l()* functions; additional

37700 searches of locale names without *.codeset* (if present), without *_territory* (if present), and without

37701 *@modifier* (if present) may be performed.

37702 For the *LANGUAGE* search, the value of the *LANGUAGE* environment variable shall be a list of

37703 one or more locale names separated by a <colon> (' : ') character. Each locale name shall be

37704 tried in the specified order. If a messages object for the locale does not exist, or cannot be

37705 opened, or is unsuitable for implementation-defined reasons (such as security), the next locale

37706 name (if any) shall be tried. If:

- 37707 • a messages object for the locale can be opened but cannot be processed without error, or
- 37708 • the messages object does not contain a string identified by *msgid*, or *msgid* and *n* for the
- 37709 *ngettext* functions,

37710 it is unspecified whether the next locale name (if any) is tried. In all other cases, the messages

37711 object for the locale shall be used.

37712 For each locale name in *LANGUAGE*, or if *LANGUAGE* is not set or is empty, or no suitable

37713 messages object is found in processing *LANGUAGE*, the pathname used to locate the messages

37714 object shall be *dirname/localename/categoryname/textdomainname.mo*, where:

- 37715 • The *dirname* part is the *dirname* argument of the most recent successful call to
37716 *bindtextdomain()* that had *textdomainname* as the *domainname* argument; any trailing <slash>
37717 characters in *dirname* shall be discarded. If a successful call to *bindtextdomain()* has not
37718 been made for *textdomainname*, an implementation-defined default directory shall be used.
- 37719 • For the *LANGUAGE* search, the *localename* part is each locale name from *LANGUAGE* in
37720 turn; if a locale name has the format *language[_territory][.codeset][@modifier]*, additional
37721 searches of locale names without *.codeset* (if present), without *_territory* (if present), and
37722 without *@modifier* (if present) may be performed; if *.codeset* is not present, additional
37723 searches of locale names with an added *.codeset* may be performed. For the single-locale
37724 search, the *localename* part is the name of the current locale, or the locale specified in an
37725 **_l()* function call, for the category named by *categoryname*. Spellings of codeset names are
37726 not standardized, and implementations may attempt to use different commonly known
37727 spellings, for example "utf8" and "UTF-8".
- 37728 • The *categoryname* part is the string "LC_MESSAGES" if *gettext()*, *dgettext()*, *ngettext()*, or
37729 *dngettext()* is called, or the locale category name corresponding to the *category* argument to
37730 *dcgettext()* or *dcngettext()*. Likewise for the **_l()* variants of all these functions.
- 37731 • For *gettext()*, *gettext_l()*, *ngettext()*, and *ngettext_l()*, the *textdomainname* part is the text
37732 domain set by the last successful call to *textdomain()*. For *dgettext()*, *dcgettext()*,
37733 *dngettext()*, *dcngettext()*, and the **_l()* variants of these functions, *textdomainname* is the text
37734 domain specified by the *domainname* argument. The *domainname* argument shall be
37735 equivalent in syntax and meaning to the *domainname* argument to *textdomain()*, except that
37736 the selection of the text domain shall affect only the *dgettext()*, *dcgettext()*, *dngettext()*, and
37737 *dcngettext()* function calls and their **_l()* variants. If the *domainname* argument is a null
37738 pointer, the text domain set by the last successful call to *textdomain()* shall be used. For all
37739 of these functions, if a successful call to *textdomain()* has not been made the default text
37740 domain "messages" shall be used.

37741 Resolution of the messages object pathname shall be performed the first time one of the *gettext*
37742 family of functions is called for a given combination of *dirname*, *localename*, *categoryname*, and
37743 *textdomainname*. It is unspecified whether the pathname is re-resolved if the combination has
37744 been used before in a call to one of the *gettext* family of functions. If *bindtextdomain()* performs
37745 pathname resolution of its *dirname* argument, only the part of the messages object pathname
37746 after *dirname* shall be resolved by the *gettext* family of functions.

37747 When one of the *gettext* family of functions returns a message string that was found in a
37748 messages object, it shall convert the codeset of the message string to the output codeset if a
37749 codeset is specified in the messages object (see *msgfmt*) and the output codeset is not the same as
37750 that codeset. If a successful call to *bind_textdomain_codeset()* has been made with the text domain
37751 of the messages object as the *domainname* argument and a non-null *codeset* argument, the output
37752 codeset shall be the *codeset* argument from the most recent such call. Otherwise, the output
37753 codeset shall be the codeset of characters in the current locale, or the provided locale object if
37754 calling one of the **_l()* functions, as specified by the *LC_CTYPE* category of the locale. The
37755 conversion shall be performed as if by a call to *iconv()* using a conversion descriptor returned by
37756 *iconv_open(<output codeset>, <messages object codeset>)*, except that if the return value of *iconv()*
37757 would be greater than zero, the non-identical conversions performed by the *gettext* family of
37758 functions need not be the same as those that such an *iconv()* call would perform. If an error
37759 prevents the codeset conversion from being performed, the *gettext* family of functions shall
37760 behave as if no message string was found in the messages object. If at least one non-identical
37761 conversion is performed that results in a fallback character (one that does not provide any
37762 information about the character it was converted from, for example, a <question-mark> or
37763 ``replacement-character''), the *gettext* family of functions may behave as if no message string was

37764 found in the messages object.

37765 RETURN VALUE

37766 The `gettext()`, `gettext_l()`, `dgettext()`, `dgettext_l()`, `dcgettext()`, and `dcgettext_l()` functions shall
37767 return the message string described in DESCRIPTION if successful. Otherwise, they shall return
37768 `msgid`.

37769 The `ngettext()`, `ngettext_l()`, `dngettext()`, `dngettext_l()`, `dcngettext()`, and `dcngettext_l()` functions
37770 shall return the message string described in DESCRIPTION if successful. Otherwise, `msgid` shall
37771 be returned if `n` is equal to 1, or `msgid_plural` if `n` is not equal to 1.

37772 The application shall ensure that it does not modify the returned string. A subsequent call to a
37773 `gettext` family function shall not overwrite or invalidate the returned string. The returned string
37774 may be invalidated by a subsequent call to `bind_textdomain_codeset()`, `bindtextdomain()`,
37775 `setlocale()`, or `textdomain()` in the same process, except for calls that only query values. The
37776 returned string shall not be invalidated by a subsequent call to `uselocale()`.

37777 ERRORS

37778 The `gettext` family of functions shall not modify `errno`. If an error occurs these functions shall
37779 return a string as described in RETURN VALUE.

37780 EXAMPLES

37781 The example code below assumes the following:

- 37782 • The implementation-defined default directory is `/system/gettextlib`.
- 37783 • The following locales are available on the target system: `en_US`, `en_GB`, `de_DE`. The
37784 codeset used for all of these locales is UTF-8.
- 37785 • The `en_AU` locale is not available on the target system.
- 37786 • The target system supports conversion from ISO/IEC 8859-1 to UTF-8.
- 37787 • The codeset used for the POSIX locale is ASCII.
- 37788 • The target system does not support conversion from ISO/IEC 8859-1 to ASCII.

37789 Furthermore, the following `.mo` files (and only the following `.mo` files) are installed:

- 37790 • `/system/gettextlib/en_US/LC_MESSAGES/mail.mo`
- 37791 • `/messagecatalogs/example/en_US/LC_MESSAGES/mail.mo`

37792 These are compiled from a portable messages object source file (dot-po file) with the following
37793 ISO/IEC 8859-1 encoded contents (see the EXTENDED DESCRIPTION of the `msgfmt` utility for a
37794 description of the dot-po file format):

```
37795 msgid ""
37796 msgstr ""
37797 "Content-Type: text/plain; charset=ISO_8859-1\n"
37798 "Plural-Forms: nplurals=4; plural= n==1?0: (n>1&& n<10)?1: (n==0)?2:3;\n"
37799 msgid "recipient"
37800 msgid_plural "recipients"
37801 msgstr[0] "1 recipient"
37802 msgstr[1] "2 to 9 recipients"
37803 msgstr[2] "no recipients"
37804 msgstr[3] "more than 9 recipients"
```

37805 `/system/gettextlib/de_DE/LC_MESSAGES/mail.mo` is compiled from a dot-po file with the
37806 following ISO/IEC 8859-1 encoded contents:

```

37807     msgid ""
37808     msgstr ""
37809     "Content-Type: text/plain; charset=ISO_8859-1\n"
37810     "Plural-Forms: nplurals=4; plural= n==1?0: (n>1&& n<5)?1: (n==0)?2:3;\n"
37811     msgid "recipient"
37812     msgid_plural "recipients"
37813     msgstr[0] "1 Empfänger"
37814     msgstr[1] "2 bis 4 Empfänger"
37815     msgstr[2] "keine Empfänger"
37816     msgstr[3] "mehr als 4 Empfänger"

37817     /messagecatalogs/example/en_GB/LC_MESSAGES/mail.mo is compiled from a dot-po file
37818     with the following ISO/IEC 8859-1 encoded contents:

37819     msgid ""
37820     msgstr ""
37821     "Content-Type: text/plain; charset=ISO_8859-1\n"
37822     "Plural-Forms: nplurals=4; plural= n==1?0: (n>1&& n<5)?1: (n==0)?2:3;\n"
37823     msgid "recipient"
37824     msgid_plural "recipients"
37825     msgstr[0] "1 recipient"
37826     msgstr[1] "2 to 4 recipients"
37827     msgstr[2] "no recipients"
37828     msgstr[3] "5 or more recipients"

37829     /messagecatalogs/example2/en_US/LC_MESSAGES/othermail.mo is not a suitable messages
37830     object file or is a suitable messages object file that does not contain the msgid "recipient".

37831     The following example demonstrates the interactions between bindtextdomain(),
37832     bind_textdomain_codeset(), textdomain(), and the gettext family of functions.

37833     unsigned long n_recipients;
37834     // strdup() is used to prevent default_domain from being invalidated by
37835     // a future call to bindtextdomain()
37836     const char *default_domain = strdup(bindtextdomain("mail", NULL));
37837     setlocale(LC_MESSAGES, "POSIX");
37838     setlocale(LC_CTYPE, "POSIX");

37839     n_recipients = 1;
37840     // The following outputs "recipient" with the same encoding as the
37841     // "recipient" argument to ngettext():
37842     printf("%s\n", ngettext("recipient", "recipients", n_recipients));

37843     n_recipients = 3;
37844     // The following outputs "recipients" with the same encoding as the
37845     // "recipients" argument to ngettext():
37846     printf("%s\n", ngettext("recipient", "recipients", n_recipients));

37847     setlocale(LC_MESSAGES, "en_US");
37848     setlocale(LC_CTYPE, "en_US");
37849     textdomain("mail");

37850     n_recipients = 1;
37851     // The following outputs "1 recipient", encoded in UTF-8:
37852     printf("%s\n", ngettext("recipient", "recipients", n_recipients));

37853     n_recipients = 3;

```

```
37854 // The following outputs "2 to 9 recipients", encoded in UTF-8:
37855 printf("%s\n", ngettext("recipient", "recipients", n_recipients));

37856 setlocale(LC_MESSAGES, "en_GB");
37857 setlocale(LC_CTYPE, "en_GB");
37858 bindtextdomain("mail", "/messagecatalogs/example/");

37859 n_recipients = 3;
37860 // The following outputs "2 to 4 recipients", encoded in UTF-8:
37861 printf("%s\n", ngettext("recipient", "recipients", n_recipients));

37862 setlocale(LC_MESSAGES, "en_US");
37863 setlocale(LC_CTYPE, "en_US");
37864 textdomain("othermail");
37865 bindtextdomain("othermail", "/messagecatalogs/example2/");

37866 n_recipients = 3;
37867 // The following outputs "recipients" with the same encoding as the
37868 // "recipients" argument to ngettext():
37869 printf("%s\n", ngettext("recipient", "recipients", n_recipients));

37870 // Because there is no locale named en_AU on the system, en_US is used:
37871 setenv("LANGUAGE", "en_AU:en_US:en_GB", 1);
37872 setlocale(LC_MESSAGES, "");
37873 setlocale(LC_CTYPE, "");
37874 bindtextdomain("mail", default_domain);

37875 // The following outputs "2 to 9 recipients", encoded in UTF-8:
37876 printf("%s\n", dngettext("mail", "recipient", "recipients", 3));

37877 textdomain("mail");
37878 bind_textdomain_codeset("mail", "UTF-8");
37879 setlocale(LC_MESSAGES, "de_DE");
37880 setlocale(LC_CTYPE, "de_DE");
37881 // Clear the LANGUAGE environment variable, otherwise it would take
37882 // precedence over the locale set above, and en_US would continue to
37883 // be used.
37884 setenv("LANGUAGE", "", 1);

37885 n_recipients = 1;
37886 // The following outputs "1 Empfänger", encoded in UTF-8:
37887 printf("%s\n", ngettext("recipient", "recipients", n_recipients));

37888 bind_textdomain_codeset("mail", "ASCII");
37889 setlocale(LC_CTYPE, "POSIX");

37890 n_recipients = 1;
37891 // The following outputs "recipient" with the same encoding as the
37892 // "recipient" argument to ngettext() - remember, the system is assumed
37893 // to not support conversion from ISO/IEC 8859-1 to ASCII:
37894 printf("%s\n", ngettext("recipient", "recipients", n_recipients));
37895 free(default_domain);
```

APPLICATION USAGE

These functions do not impose a limit on message length. Note that translated strings typically have a different length than the input strings, possibly much longer, and applications using these translations in formatted text (for example, aligned columns for a table) should take that into account.

The `dcgettext()`, `dcgettext_l()`, `dcngettext()`, and `dcngettext_l()` functions are useful to retrieve locale-specific strings for a category other than `LC_MESSAGES`. For example, they can be used to obtain a time format string from the `LC_TIME` category; because the locale setting of `LC_TIME` and `LC_MESSAGES` can be different, using the other `gettext` family functions in such a case might cause an undesired result. All of the functions in the `gettext` family of functions, except `dcgettext()`, `dcgettext_l()`, `dcngettext()`, and `dcngettext_l()`, search for messages objects only in the `LC_MESSAGES` category.

Implementations typically, but are not required to, `mmap()` the messages object file the first time one of the `gettext` family of functions is called, and keep that map in place until it is no longer expected to be used. For example, a successful call to `bindtextdomain()` will typically cause the next call to one of the `gettext` family of functions to `munmap()` the previous file and `mmap()` the new file. Applications should not rely on this behavior, however: the implementation is allowed to cache previously used maps, or not use `mmap()` at all and reopen the file each time one of the `gettext` family of functions is called.

The `msgid` and `msgid_plural` arguments are typically in (US) English. The arguments are always used in the POSIX or C locale, and when a `gettext` family function encounters an error, so they should not be abstract message identifiers (for example, "message 123") and they should only use characters in the portable character set (to avoid outputting byte sequences that are not valid characters in the current output codeset). If the `xgettext` utility is used to extract the `msgid` and `msgid_plural` arguments from C source files into a template dot-po file, the arguments must be string literals in order for the resulting file to be useful to translators.

The strings returned by the `gettext` family of functions are not guaranteed to contain only characters that are valid in the current output codeset. In particular, byte sequences that do not form valid characters can occur when:

- The `msgid` or `msgid_plural` arguments use characters outside the portable character set.
- The messages object file does not specify a character set and uses characters outside the portable character set.

The strings returned by the `gettext` family of functions are guaranteed to remain valid until invalidated as described in the RETURN VALUE section. This includes strings that are created by codeset conversion; those strings are freed by the implementation, not the application. Thus, it is safe to call `gettext` family functions multiple times in situations such as:

```
printf("%s %s\n", gettext("foo"), gettext("bar"));
```

RATIONALE

Although the return type of these functions ought to be `const char *`, it is `char *` to match historical practice.

The `gettext` family of functions is frequently used in reporting errors. In fact, it is possible to have an application that attempts to create an error message that combines a translated string via `gettext()` with an error string provided by `strerror()`. The standard requires that the `gettext` family of functions does not modify `errno`, so that an application need not worry about complications of providing sequencing points to capture a stable value of `errno` prior to the translation of the error message, and so that the user will still get a somewhat useful string (even if it is the untranslated original string) on any failure.

37943 There are no wide character equivalents for these functions; historically no implementation is
37944 known to exist, and the multi-byte message returned from these functions can, in most instances,
37945 be converted to wide characters by the application if desired.

37946 Some historical *gettext* implementations returned the translated string from the messages object
37947 without codeset conversion if *iconv_open()* fails. This is considered to be a bug in those
37948 implementations.

37949 **FUTURE DIRECTIONS**

37950 None.

37951 **SEE ALSO**

37952 *bindtextdomain()*, *catopen()*, *iconv()*, *setlocale()*, *uselocale()*

37953 XBD [<libintl.h>](#), [<limits.h>](#)

37954 XCU *gettext*, *msgfmt*, *xgettext*

37955 **CHANGE HISTORY**

37956 First released in Issue 8.

37957

47012 **NAME**

47013 ngettext, ngettext_l — message handling functions

47014 **SYNOPSIS**

47015 #include <libintl.h>

47016 char *ngettext(const char *msgid, const char *msgid_plural,
47017 unsigned long int n);47018 char *ngettext_l(const char *msgid, const char *msgid_plural,
47019 unsigned long int n, locale_t locale);47020 **DESCRIPTION**47021 Refer to *gettext*.

61783 **NAME**

61784 setresgid — set real group ID, effective group ID, and saved set-group-ID

61785 **SYNOPSIS**

61786 XSI #include <unistd.h>

61787 int setresgid(gid_t rgid, gid_t egid, gid_t sgid);

61788

61789 **DESCRIPTION**61790 The *setresgid()* function shall set the real group ID, effective group ID, and saved set-group-ID of
61791 the calling process to the values specified by *rgid*, *egid*, and *sgid*, respectively.61792 If an argument is -1 , the corresponding ID shall not be changed.61793 Only a process with appropriate privileges can set the real group ID, effective group ID, and
61794 saved set-group-ID to any valid value.61795 A non-privileged process can set its real group ID, effective group ID, and saved set-group-ID,
61796 each to one of the values that it currently holds in its real group ID, effective group ID, or saved
61797 set-group-ID.61798 The real group ID, effective group ID, and saved set-group-ID can be set to different values in
61799 the same call.

61800 Any supplementary group IDs of the calling process shall remain unchanged.

61801 **RETURN VALUE**61802 Upon successful completion, 0 shall be returned. Otherwise, -1 shall be returned and *errno* set to
61803 indicate the error, and none of the IDs shall be changed.61804 **ERRORS**61805 The *setresgid()* function shall fail if:61806 [EINVAL] The value of the *rgid*, *egid*, or *sgid* argument is invalid or out-of-range.61807 [EPERM] The calling process does not have appropriate privileges and an attempt was
61808 made to change the real group ID, effective group ID, or saved set-group-ID to
61809 a value that is not currently present in one of those IDs.61810 **EXAMPLES**

61811 None.

61812 **APPLICATION USAGE**

61813 None.

61814 **RATIONALE**

61815 None.

61816 **FUTURE DIRECTIONS**

61817 None.

61818 **SEE ALSO**61819 *exec*, *getegid()*, *geteuid()*, *getgid()*, *getresgid()*, *getresuid()*, *getuid()*, *setegid()*, *seteuid()*, *setgid()*,
61820 *setregid()*, *setresuid()*, *setreuid()*, *setuid()*

61821 XBD <unistd.h>

61822 **CHANGE HISTORY**
61823 First released in Issue 8.
61824

NAME

61825 setresuid — set real user ID, effective user ID, and saved set-user-ID
 61826

SYNOPSIS

```
61828 XSI #include <unistd.h>
61829 int setresuid(uid_t ruid, uid_t euid, uid_t suid);
61830
```

DESCRIPTION

61831 The *setresuid()* function shall set the real user ID, effective user ID, and saved set-user-ID of the
 61832 calling process to the values specified by *ruid*, *euid*, and *suid*, respectively.
 61833

61834 If an argument is -1 , the corresponding ID shall not be changed.

61835 Only a process with appropriate privileges can set the real user ID, effective user ID, and saved
 61836 set-user-ID to any valid value.

61837 A non-privileged process can set its real user ID, effective user ID, and saved set-user-ID, each to
 61838 one of the values that it currently holds in its real user ID, effective user ID, or saved set-user-ID.

61839 The real user ID, effective user ID, and saved set-user-ID can be set to different values in the
 61840 same call.

RETURN VALUE

61841 Upon successful completion, 0 shall be returned. Otherwise, -1 shall be returned and *errno* set to
 61842 indicate the error, and none of the IDs shall be changed.
 61843

ERRORS

61844 The *setresuid()* function shall fail if:

61846 [EINVAL] The value of the *ruid*, *euid*, or *suid* argument is invalid or out-of-range.

61847 [EPERM] The calling process does not have appropriate privileges and an attempt was
 61848 made to change the real user ID, effective user ID, or saved set-user-ID to a
 61849 value that is not currently present in one of those IDs or an attempt was made
 61850 to change the real user ID to a value not permitted by the implementation.

EXAMPLES

61851 None.
 61852

APPLICATION USAGE

61853 None.
 61854

RATIONALE

61855 None.
 61856

FUTURE DIRECTIONS

61857 None.
 61858

SEE ALSO

61860 *exec*, *getegid()*, *geteuid()*, *getgid()*, *getresgid()*, *getresuid()*, *getuid()*, *setegid()*, *seteuid()*, *setgid()*,
 61861 *setregid()*, *setresgid()*, *setreuid()*, *setuid()*

61862 XBD <unistd.h>

CHANGE HISTORY

61863 First released in Issue 8.
 61864
 61865

69197 **NAME**

69198 textdomain — text domain manipulation function |

69199 **SYNOPSIS**

69200 #include <libintl.h> |

69201 char *textdomain(const char *domainname); |

69202 **DESCRIPTION**69203 Refer to *bindtextdomain()*. |

84828 **STDOUT**84829 When the `-v` option is specified, standard output shall be formatted as:84830 `"%s\n", <pathname or command>`84831 When the `-V` option is specified, standard output shall be formatted as:84832 `"%s\n", <unspecified>`84833 **STDERR**

84834 The standard error shall be used only for diagnostic messages.

84835 **OUTPUT FILES**

84836 None.

84837 **EXTENDED DESCRIPTION**

84838 None.

84839 **EXIT STATUS**84840 When the `-v` or `-V` options are specified, the following exit values shall be returned:

84841 0 Successful completion.

84842 >0 The *command_name* could not be found or an error occurred.

84843 Otherwise, the following exit values shall be returned:

84844 126 The utility specified by *command_name* was found but could not be invoked.84845 127 An error occurred in the *command* utility or the utility specified by *command_name* could not
84846 be found.84847 Otherwise, the exit status of *command* shall be that of the simple command specified by the
84848 arguments to *command*.84849 **CONSEQUENCES OF ERRORS**

84850 Default.

84851 **APPLICATION USAGE**84852 This utility is required to be intrinsic. See [Section 1.7](#) (on page 2336) for details.84853 The order for command search allows functions to override regular built-ins and path searches.
84854 This utility is necessary to allow functions that have the same name as a utility to call the utility
84855 (instead of a recursive call to the function).84856 The system default path is available using *getconf*; however, since *getconf* may need to have the
84857 *PATH* set up before it can be called itself, the following can be used:84858 `command -p getconf PATH`84859 There are some advantages to suppressing the special characteristics of special built-ins on
84860 occasion. For example:84861 `command exec > unwritable-file`84862 does not cause a non-interactive script to abort, so that the output status can be checked by the
84863 script.84864 The *command*, *env*, *nohup*, *time*, *timeout*, and *xargs* utilities have been specified to use exit code 127 +
84865 if an error occurs so that applications can distinguish “failure to find a utility” from “invoked
84866 utility exited with an error indication”. The value 127 was chosen because it is not commonly
84867 used for other meanings; most utilities use small values for “normal error conditions” and the
84868 values above 128 can be confused with termination due to receipt of a signal. The value 126 was

88560 XSI **NLSPATH** Determine the location of message catalogs for the processing of *LC_MESSAGES*.
 88561 **PATH** Determine the location of the *utility*, as described in XBD Chapter 8 (on page 155).
 88562 If *PATH* is specified as a *name=value* operand to *env*, the *value* given shall be used in
 88563 the search for *utility*.

88564 **ASYNCHRONOUS EVENTS**

88565 Default.

88566 **STDOUT**

88567 If no *utility* operand is specified, each *name=value* pair in the resulting environment shall be
 88568 written in the form:

88569 "*%s=%s\n*", *<name>*, *<value>*

88570 If the *utility* operand is specified, the *env* utility shall not write to standard output.

88571 **STDERR**

88572 The standard error shall be used only for diagnostic messages.

88573 **OUTPUT FILES**

88574 None.

88575 **EXTENDED DESCRIPTION**

88576 None.

88577 **EXIT STATUS**

88578 If *utility* is invoked, the exit status of *env* shall be the exit status of *utility*; otherwise, the *env*
 88579 utility shall exit with one of the following values:

88580 0 The *env* utility completed successfully.

88581 1–125 An error occurred in the *env* utility.

88582 126 The utility specified by *utility* was found but could not be invoked.

88583 127 The utility specified by *utility* could not be found.

88584 **CONSEQUENCES OF ERRORS**

88585 Default.

88586 **APPLICATION USAGE**

88587 The *command*, *env*, *nice*, *nohup*, *time*, *timeout*, and *xargs* utilities have been specified to use exit +
 88588 code 127 if an error occurs so that applications can distinguish “failure to find a utility” from
 88589 “invoked utility exited with an error indication”. The value 127 was chosen because it is not
 88590 commonly used for other meanings; most utilities use small values for “normal error
 88591 conditions” and the values above 128 can be confused with termination due to receipt of a
 88592 signal. The value 126 was chosen in a similar manner to indicate that the utility could be found,
 88593 but not invoked. Some scripts produce meaningful error messages differentiating the 126 and
 88594 127 cases. The distinction between exit codes 126 and 127 is based on KornShell practice that
 88595 uses 127 when all attempts to *exec* the utility fail with [ENOENT], and uses 126 when any
 88596 attempt to *exec* the utility fails for any other reason.

88597 Historical implementations of the *env* utility use the *execvp()* or *execlp()* functions defined in the
 88598 System Interfaces volume of POSIX.1-202x to invoke the specified utility; this provides better
 88599 performance and keeps users from having to escape characters with special meaning to the shell.
 88600 Therefore, shell functions, special built-ins, and built-ins that are only provided by the shell are
 88601 not found by this type of *env* implementation. However, *env* can be implemented as a shell built-
 88602 in, in which case it may be able to execute shell functions and built-ins. An application wishing
 88603 to ensure execution of a non-built-in utility can use *exec* in a subshell for this purpose.

94098 **NAME**

94099 gettext, ngettext — retrieve text string from messages object +

94100 **SYNOPSIS** +

94101 gettext [-e|-E] [-d *textdomain*] [*textdomain*] *msgid* +

94102 gettext [-e|-E] [-n] -s [-d *textdomain*] *msgid*... +

94103 ngettext [-e|-E] [-d *textdomain*] [*textdomain*] *msgid msgid_plural n* +

94104 **DESCRIPTION** +

94105 The *gettext* and *ngettext* utilities shall write to standard output the message string(s) that would +
 94106 result from the following calls to functions defined in the System Interfaces volume of +
 94107 POSIX.1-202x: +

```

94108            if (textdomainname == NULL || textdomainname[0] == '\0') +
94109                message_string = msgid; +
94110            else { +
94111                setlocale(LC_ALL, ""); +
94112                if (textdomaindir != NULL) +
94113                    bindtextdomain(textdomainname, textdomaindir); +
94114                if (msgid_plural == NULL) +
94115                    message_string = dgettext(textdomainname, msgid); +
94116                else +
94117                    message_string = dngettext(textdomainname, msgid, msgid_plural, n); +
94118            } +
  
```

94119 where: +

94120 • The *textdomaindir* variable is a string containing the value of the *TEXTDOMAINDIR* +
 94121 environment variable, if set and not empty, or is NULL otherwise. +

94122 • The *textdomainname* variable is a string containing the text domain name obtained from, in +
 94123 decreasing order of precedence: +

94124 — The optional operand *textdomain*, if present +

94125 — The *-d textdomain* option, if specified +

94126 — The *TEXTDOMAIN* environment variable, if set and not empty +

94127 If the text domain name cannot be obtained from these sources, the *textdomainname* +
 94128 variable is NULL. +

94129 • If the *-s* option of *gettext* is not specified and for the *ngettext* utility, the *msgid* variable is a +
 94130 string containing: +

94131 — The value of the *msgid* operand, if the *-E* option is specified +

94132 — The value of the *msgid* operand with C-language escape sequences processed (see +
 94133 below), if the *-e* option is specified +

94134 — The value of the *msgid* operand with C-language escape sequences optionally +
 94135 processed (see below), otherwise +

94136 • If the *-s* option of *gettext* is specified, the *msgid* variable is a string containing: +

94137 — The value of each *msgid* operand in turn, if the *-E* option is specified or neither the *-e* +
 94138 nor the *-E* option is specified +

- 94139 — The value of each *msgid* operand in turn with C-language escape sequences +
 94140 processed (see below), if the **-e** option is specified +
- 94141 • For the *gettext* utility, the *msgid_plural* variable is NULL. For the *ngettext* utility, the +
 94142 *msgid_plural* variable is a string containing: +
 - 94143 — The value of the *msgid_plural* operand, if the **-E** option is specified +
 - 94144 — The value of the *msgid_plural* operand with C-language escape sequences processed +
 94145 (see below), if the **-e** option is specified +
 - 94146 — The value of the *msgid_plural* operand with C-language escape sequences optionally +
 94147 processed (see below), otherwise +
 - 94148 • For the *gettext* utility, the *n* variable is 1 (one). For the *ngettext* utility the *n* variable is the *n* +
 94149 operand, parsed as an integer as if by using the *strtoul()* function with a *base* argument of +
 94150 10. +
- 94151 When C-language escape sequences are processed, they shall be processed as specified for +
 94152 character string literals in the ISO C standard, except that *universal-character-name* escape +
 94153 sequences need not be supported. Implementations may also support a <backslash> 'c' escape +
 94154 sequence; if supported, the '\c' and all characters following it shall be removed and, if the **-s** +
 94155 option is specified, the behavior shall be as if the **-n** option is also specified. +
- 94156 For the *ngettext* utility, and for the *gettext* utility if the **-s** option is not specified, the resulting +
 94157 message string shall be written to standard output. If the **-s** option of *gettext* is specified, the +
 94158 resulting message string for each *msgid* shall be written to standard output with consecutive +
 94159 message strings separated by a single <space> character and, if the **-n** option is not specified, a +
 94160 <newline> shall be written after the last message string. If the **-s** and **-n** options are specified, +
 94161 the trailing <newline> shall be omitted. +
- 94162 Under conditions where the *textdomainname* variable in the above code would be NULL, these +
 94163 utilities may write a diagnostic message to standard error and exit with non-zero status. +
- 94164 **OPTIONS** +
- 94165 These utilities shall conform to XBD [Section 12.2](#) (on page 201). +
- 94166 The following options shall be supported: +
- 94167 **-d** *textdomain* +
 94168 Retrieve the translated message from the domain *textdomain*, if *textdomain* is not +
 94169 specified as an operand. +
 - 94170 **-e** Process C-language escape sequences in *msgid* and *msgid_plural* operands. +
 - 94171 **-E** Do not process C-language escape sequences in *msgid* and *msgid_plural* operands. +
- 94172 The *gettext* utility shall also support the following options: +
- 94173 **-n** Modify the behavior of the **-s** option such that a <newline> is not appended to the +
 94174 output. +
 - 94175 **-s** Separate the message strings obtained from each *msgid* operand with <space> +
 94176 characters in the output, and (if **-n** is not also specified) append a <newline> to the +
 94177 output. +
- 94178 If neither of the mutually exclusive **-e** and **-E** options is specified, it is unspecified which is the +
 94179 default, except that if the **-s** option of *gettext* is specified then **-E** shall be the default. +

94180	OPERANDS		+
94181		The following operands shall be supported:	+
94182	<i>textdomain</i>	A text domain name used to retrieve the translated message. This shall override	+
94183		the specification by the <code>-d</code> option, if present.	+
94184	<i>msgid</i>	A key to retrieve the translated message.	+
94185	<i>msgid_plural</i>	A default plural if no corresponding plural message can be found.	+
94186	<i>n</i>	A non-negative decimal integer to be used as the <i>n</i> argument to <code>dngettext()</code> (see the	+
94187		DESCRIPTION).	+
94188	STDIN		+
94189		Not used.	+
94190	INPUT FILES		+
94191		The input files are messages object files (see <i>msgfmt</i>).	+
94192	ENVIRONMENT VARIABLES		+
94193		The following environment variables shall affect the execution of <i>gettext</i> and <i>ngettext</i> :	+
94194	<i>LANG</i>	Provide a default value for the internationalization variables that are unset or null.	+
94195		(See XBD Section 8.2 (on page 157) for the precedence of internationalization	+
94196		variables used to determine the values of locale categories.)	+
94197	XSI	<i>LANGUAGE</i> Determine the location of messages objects if <i>NLSPATH</i> is not set or the evaluation	+
94198		of <i>NLSPATH</i> did not lead to a suitable messages object being found.	+
94199		<i>LC_ALL</i> If set to a non-empty string value, override the values of all the other	+
94200		internationalization variables.	+
94201		<i>LC_MESSAGES</i>	+
94202		Determine the locale name used to locate messages objects, and the locale that	+
94203		should be used to affect the format and contents of diagnostic messages written to	+
94204		standard error.	+
94205	XSI	<i>NLSPATH</i> Determine the location of messages objects and message catalogs.	+
94206		<i>TEXTDOMAIN</i>	+
94207		Specify the text domain name. (See XBD Section 3.374 (on page 81).)	+
94208		<i>TEXTDOMAINDIR</i>	+
94209	XSI	Specify the pathname to the messages object hierarchy. <i>NLSPATH</i> shall have	+
94210		precedence over <i>TEXTDOMAINDIR</i> .	+
94211	ASYNCHRONOUS EVENTS		+
94212		Default.	+
94213	STDOUT		+
94214		See DESCRIPTION.	+
94215	STDERR		+
94216		The standard error shall be used only for diagnostic messages.	+
94217	OUTPUT FILES		+
94218		None.	+

94219	EXTENDED DESCRIPTION	+
94220	None.	+
94221	EXIT STATUS	+
94222	The following exit values shall be returned:	+
94223	0 Successful completion.	+
94224	>0 An error occurred.	+
94225	CONSEQUENCES OF ERRORS	+
94226	Default.	+
94227	APPLICATION USAGE	+
94228	Since it is unspecified which of the <code>-e</code> or <code>-E</code> options is the default, except when the <code>-s</code> option of <code>gettext</code> is specified, portable applications need to ensure that <code>-e</code> , <code>-E</code> , or (for <code>gettext</code>) <code>-s</code> is specified	+
94229	whenever a <code>msgid</code> or <code>msgid_plural</code> operand contains, or might contain, a <code><backslash></code> character.	+
94230		+
94231	Note that, unless the <code>-s</code> option of <code>gettext</code> is specified without <code>-n</code> , the message(s) written to	+
94232	standard output are not followed by a <code><newline></code> . (Therefore the output only ends with a	+
94233	<code><newline></code> if the last message ends with one.)	+
94234	Both <code>msgid</code> and <code>msgid_plural</code> should be properly quoted for the shell.	+
94235	EXAMPLES	+
94236	The following examples assume that the following portable messages object source file (<code>dot-po</code>	+
94237	file) has been compiled to a valid file mail.mo by the <code>msgfmt</code> utility. See the EXTENDED	+
94238	DESCRIPTION section of the <code>msgfmt</code> utility for a description of the <code>dot-po</code> file format.	+
94239	<code>msgid ""</code>	+
94240	<code>msgstr ""</code>	+
94241	<code>"Content-Type: text/plain; charset=utf-8\n"</code>	+
94242	<code>"Plural-Forms: nplurals=4; plural=n==1?0: (n>1&&n<=10)?1: (n==0)?2:3;\n"</code>	+
94243	<code>msgid "recipient"</code>	+
94244	<code>msgid_plural "recipients"</code>	+
94245	<code>msgstr[0] "1 recipient"</code>	+
94246	<code>msgstr[1] "2 to 10 recipients"</code>	+
94247	<code>msgstr[2] "no recipients"</code>	+
94248	<code>msgstr[3] "more than 10 recipients"</code>	+
94249	<code>msgid "%d attachment\n"</code>	+
94250	<code>msgid_plural "%d attachments\n"</code>	+
94251	<code>msgstr[0] "1 (%d) attachment\n"</code>	+
94252	<code>msgstr[1] "2 to 10 (%d) attachments\n"</code>	+
94253	<code>msgstr[2] "no (%d) attachments\n"</code>	+
94254	<code>msgstr[3] "more than 10 (%d) attachments\n"</code>	+
94255	They also assume that mail.mo is installed in the directory that <code>gettext</code> and <code>ngettext</code> search for the	+
94256	current locale. See the OPTIONS and ENVIRONMENT VARIABLES sections above and the	+
94257	description of <code>gettext()</code> for details on how this search is performed.	+
94258	The command	+
94259	<code>ngettext -d mail recipient recipients 0</code>	+
94260	will write "no recipients".	+
94261	The command	+

94262 `ngettext -d mail recipient recipients 1` +
94263 will write "1 recipient". +
94264 The command +
94265 `ngettext -d mail recipient recipients 5` +
94266 will write "2 to 10 recipients". +
94267 The command +
94268 `ngettext -d mail recipient recipients 11` +
94269 will write "more than 10 recipients". +
94270 The command +
94271 `ngettext -d mail Call Calls 1` +
94272 will write "Call". Note that "Call" is not in the messages object. +
94273 The command +
94274 `ngettext -d mail Call Calls 0` +
94275 will write "Calls". +
94276 The command +
94277 `ngettext -d mail Call Calls 10` +
94278 will write "Calls". +
94279 The command +
94280 `ngettext -ed mail "%d attachment\n" "%d attachments\n" 1` +
94281 will write the same as +
94282 `printf "1 (%d) attachment\n"` +
94283 (i.e. "1 (%d) attachment" followed by a <newline> character). The output of *ngettext* can be +
94284 used as a format string for *printf*. +
94285 The command +
94286 `printf "$ (ngettext -ed mail "%d attachment\n" "%d attachments\n" 1)" 10` +
94287 will write the same as +
94288 `printf "1 (%d) attachment\n" 10` +
94289 (i.e. "1 (10) attachment" followed by a <newline> character). +
94290 The command +
94291 `ngettext -e -d mail "\tsubject\n" "\tsubjects\n" 0` +
94292 will write the same as +
94293 `printf "\tsubjects\n"` +
94294 (i.e. a <tab> character, followed by "subjects" followed by a <newline> character). Note that +
94295 `"\tsubject\n"` is not in the messages object. +
94296 The command +

94297	<code>printf "%s\n" "\$(gettext -E -d mail "subject" "subjects" 0) "</code>	+
94298	will write the same as	+
94299	<code>printf "subjects\n"</code>	+
94300	(i.e. "subjects" followed by a <newline> character). Note that "subject" is not in the	+
94301	messages object.	+
94302	The command	+
94303	<code>gettext -s -d mail "recipient"</code>	+
94304	will write "1 recipient" followed by a <newline> character.	+
94305	The command	+
94306	<code>gettext -s -n -d mail "recipient"</code>	+
94307	will write "1 recipient" without a <newline> character.	+
94308	RATIONALE	+
94309	Historical implementations did not support the '\a' C-language escape sequence. This	+
94310	standard requires it to be supported for consistency with other utilities that support the table in	+
94311	XBD Chapter 5 (on page 101).	+
94312	Unlike other standard utilities, the behavior of <i>gettext</i> and <i>gettext</i> is not undefined when	+
94313	<i>NLSPATH</i> overrides the system default path; see XBD Section 8.2 (on page 157). This is so that	+
94314	applications can use these utilities to obtain message strings from messages objects in other	+
94315	locations. However, it also means that they need to be implemented in such a way that they do	+
94316	not do anything that would result in undefined behavior when they need to write a diagnostic	+
94317	message. In particular, they should not use a string obtained from a message catalog or a	+
94318	messages object as a format string (or should only do so after checking that the string contains	+
94319	the correct conversions).	+
94320	FUTURE DIRECTIONS	+
94321	None.	+
94322	SEE ALSO	+
94323	<i>msgfmt</i> , <i>printf</i>	+
94324	XBD Chapter 7 (on page 115), Chapter 8 (on page 155), Section 12.2 (on page 201)	+
94325	XSH <i>gettext</i> , <i>iconv()</i> , <i>setlocale()</i>	+
94326	CHANGE HISTORY	+
94327	First released in Issue 8.	+
94328		+

101620 **NAME**

101621 msgfmt — create messages objects from portable messages object source files

101622 **SYNOPSIS**

101623 msgfmt [-cfSv] [-D *dir*] [-o *outputfile*] *pathname...*

101624 **DESCRIPTION**

101625 The *msgfmt* utility shall create messages object files from portable messages object source files
101626 (dot-po files).

101627 A dot-po file contains messages to be output by system commands or by applications. The
101628 messages in these files should be able to be translated to any language supported by the system.

101629 The *msgfmt* utility shall interpret message strings for output as characters according to the
101630 codeset specified in the dot-po file or, if not present, the current setting of the *LC_CTYPE* locale
101631 category.

101632 **OPTIONS**

101633 The *msgfmt* utility shall conform to XBD [Section 12.2](#) (on page 201).

101634 The following options shall be supported:

101635 **-c** If this option and **-v** are both specified, *msgfmt* shall detect and diagnose input file
101636 abnormalities which might represent translation errors. The **msgid** and **msgstr**
101637 strings shall be compared. It shall be considered abnormal if one string starts or
101638 ends with a <newline> while the other does not. Also, if the flag **c-format** appears
101639 in a "#, " comment for a **msgid** directive (see EXTENDED DESCRIPTION), it shall
101640 be considered abnormal if the strings do not have the same number of '%'
101641 conversion specifiers, or if corresponding conversion specifiers take different
101642 argument types (see XSH *fprintf()*, on page 909). If an abnormality is detected, the
101643 exit status shall be non-zero and a diagnostic message shall be output. Additional
101644 checks beyond those described here may also be performed. These checks may
101645 produce diagnostics or informational messages and need not affect the exit status.
101646 If **-c** is specified without **-v** or **-v** is specified without **-c**, the behavior is
101647 unspecified.

101648 **-D *dir*** Add *dir* to the list of directories to search for input files.

101649 **-f** Use fuzzy entries in output. If this option is not specified, fuzzy entries shall not be
101650 included in the output.

101651 **-o *outputfile***

101652 Specify the name of an output file to be used instead of the default filename(s)
101653 specified in EXTENDED DESCRIPTION. All **domain *domainname*** directives in the
101654 dot-po file(s) shall be ignored.

101655 **-S** Append the suffix **.mo** to each generated messages object filename if it does not
101656 have this suffix.

101657 **-v** See **-c**.

101658 **OPERANDS**

101659 The following operand shall be supported:

101660 *pathname* A pathname of a dot-po file.

- 101661 **STDIN**
- 101662 Not used.
- 101663 **INPUT FILES**
- 101664 The input files shall be text files in the format described in EXTENDED DESCRIPTION.
- 101665 **ENVIRONMENT VARIABLES**
- 101666 The following environment variables shall affect the execution of *msgfmt*:
- 101667 *LANG* Provide a default value for the internationalization variables that are unset or null.
101668 (See XBD Section 8.2 (on page 157) for the precedence of internationalization
101669 variables used to determine the values of locale categories.)
- 101670 XSI *LANGUAGE* Determine the location of messages objects if *NLSPATH* is not set or the evaluation
101671 of *NLSPATH* did not lead to a suitable messages object being found.
- 101672 *LC_ALL* If set to a non-empty string value, override the values of all the other
101673 internationalization variables.
- 101674 *LC_CTYPE* Determine the locale for the interpretation of sequences of bytes of text data as
101675 characters (for example, single-byte as opposed to multi-byte characters in
101676 arguments and input files).
- 101677 *LC_MESSAGES*
- 101678 Determine the locale name used to locate messages objects, and the locale that
101679 should be used to affect the format and contents of diagnostic messages written to
101680 standard error.
- 101681 XSI *NLSPATH* Determine the location of messages objects and message catalogs.
- 101682 **ASYNCHRONOUS EVENTS**
- 101683 Default.
- 101684 **STDOUT**
- 101685 Not Used.
- 101686 **STDERR**
- 101687 The standard error shall be used for diagnostic messages and may also be used for warning
101688 messages. If the *-c* and *-v* options are specified, additional unspecified informational messages
101689 may be written to standard error.
- 101690 **OUTPUT FILES**
- 101691 The format of the created messages object files is unspecified.
- 101692 **EXTENDED DESCRIPTION**
- 101693 The *msgfmt* utility shall accept portable messages object source files (dot-po files) in the
101694 following format.
- 101695 A dot-po file contains zero or more lines, with each non-blank line containing a comment, a
101696 statement, or a statement continuation. A comment has an unquoted <number-sign> ('#') as
101697 the first non-<blank> character and ends with the next <newline> character. A statement
101698 continuation is a double-quoted string on a line by itself, optionally preceded and/or followed
101699 by <blank> characters, and the string shall be concatenated with the string on the previous
101700 statement line. If a comment occurs between a statement and a statement continuation, the
101701 behavior is unspecified. All other comments, except for comments beginning with <number-
101702 sign><comma> ("#, "), and blank lines shall be ignored.
- 101703 The format of a statement is:
- 101704 *directive value*

101705 The *directive* starts at the first non-<blank> character of the line and is separated from the *value*
 101706 by one or more <blank> characters. The *value* consists of a double-quoted string optionally
 101707 followed by <blank> characters. Zero or more statement continuation lines (see above) can
 101708 follow the statement. The following directives shall be supported:

```
101709 domain domainname
101710 msgid message_identifier
101711 msgid_plural untranslated_string_plural
101712 msgstr message_string
101713 msgstr[index] message_string
```

101714 A dot-po file consists of zero or more sections. Each section specifies the messages to be
 101715 processed in a domain. The first directive in each section shall be a **domain** directive (except for
 101716 the first section which shall behave as if

```
101717 domain "messages"
```

101718 had been specified if the first directive is not a **domain** directive).

101719 The behavior of the **domain** directive is affected by the options used. See OPTIONS for the
 101720 behavior when the **-o** option is specified. If the **-o** option is not specified, all data obtained from
 101721 the non-**domain** directives in a dot-po section shall be output to the messages object file named
 101722 *domainname.mo* when the **-S** option is specified. When the **-S** option is not specified, it is
 101723 implementation-defined whether *domainname* or *domainname.mo* is used.

101724 If multiple **domain** directives specify the same *domainname*, the sections shall be processed as if
 101725 there was only one section that starts with a **domain** *domainname* statement which contained the
 101726 statements of the sections, in the same order, excluding all but the first **domain** *domainname*
 101727 statement.

101728 Within each section, there can be a header. A header is identified by having a **msgid** directive
 101729 with the empty string ("") as the *message_identifier* immediately followed by a statement
 101730 containing a **msgstr** directive. The *message_string* in this **msgstr** statement in a header shall be
 101731 treated specially. If *message_string* contains a specification of the form:

```
101732 "nplurals=count; plural=expression"
```

101733 then *count* indicates the number of plural forms for messages in that domain, and *expression* is a
 101734 C-language expression that evaluates to an unsigned integer value which determines the
 101735 **msgstr**[*index*] directive to be used. The value of *expression* is used as the index value. The
 101736 variable *n* in *expression* is assigned the value of the *n* argument to the *ngettext()*, *ngettext_l()*,
 101737 *dngettext()*, *dngettext_l()*, *dcngettext()*, and *dcngettext_l()* functions or of the *n* operand of the
 101738 *ngettext* utility before *expression* is evaluated. The application shall ensure that *expression*
 101739 evaluates to a non-negative value less than *count* for all *n* that can be supplied by the
 101740 aforementioned functions and utility.

101741 If *message_string* in the header contains a specification of the form:

```
101742 "charset=codeset"
```

101743 then *codeset* indicates the codeset to be used to encode the message strings in this section's
 101744 domain (overriding *LC_CTYPE*). If the output string's codeset is different from the message
 101745 string's codeset, codeset conversion from the message string's codeset to the output string's
 101746 codeset shall be performed by the *gettext* family of functions and by the *gettext* and *ngettext*
 101747 utilities. See XSH *gettext* and *gettext*. The output string's codeset shall be determined by the
 101748 current or specified locale's codeset.

101749 **Note:** It is the responsibility of translators to ensure that the characters they enter into message strings
101750 in a dot-po file are encoded in the codeset specified in the header.

101751 If a header is present in a section, the application shall ensure that the header is provided by the
101752 first **msgid** directive in that section.

101753 After the header, if present, zero or more messages are identified by a **msgid** directive with a
101754 *message_identifier* that is not an empty string. Each of these directives start a subsection that is
101755 used to get a translated message from the *gettext* family of functions and from the *gettext* and
101756 *ngettext* utilities. If the *message_identifier* string is the string identified by the *gettext* family of
101757 functions *msgid* argument or by the *gettext* and *ngettext* utility *msgid* operand, this subsection
101758 specifies how that translation is to be processed.

101759 If there is only a singular form for the given *message_identifier*, the application shall ensure that
101760 the statement containing the **msgid** directive is immediately followed by a **msgstr** directive.

101761 If there are plural forms for the given *message_identifier* and the header for this section exists and
101762 contains an

101763 "*nplurals=count; plural=expression*"

101764 specification, the application shall ensure that the statement containing the **msgid** directive is
101765 immediately followed by a **msgid_plural** directive and that each statement containing a
101766 **msgid_plural** directive is followed by *count* statements containing **msgstr[index]** directives,
101767 starting with **msgstr[0]** and ending with **msgstr[count-1]** in monotonically increasing order. If a
101768 header for this section does not exist or does not contain an

101769 "*nplurals=count; plural=expression*"

101770 specification, the application shall ensure that no **msgid_plural** or **msgstr[index]** directives are
101771 used in this section.

101772 For example, if the header's *message_string* contains the specification:

101773 "*nplurals=2; plural= n == 1 ? 0 : 1*"

101774 there are two forms in the domain; **msgstr[0]** is used if *n* is equal to 1, otherwise **msgstr[1]** is
101775 used. For another example, if the header's *message_string* contains:

101776 "*nplurals=3; plural= n == 1 ? 0 : n == 2 ? 1 : 2*"

101777 there are three forms in the domain; **msgstr[0]** is used if *n* is equal to 1, **msgstr[1]** is used if *n* is
101778 equal to 2, otherwise **msgstr[2]** is used.

101779 C-language escape sequences in strings shall be processed as specified for character string
101780 literals in the ISO C standard, except that *universal-character-name* escape sequences need not be
101781 supported.

101782 Comments in a dot-po file can be in one of the following formats:

101783 # : *reference*
101784 # . *utility-added-comments*
101785 # , *flag*
101786 # *translator-comments* (where *translator-comments* does not begin with '.', ': ' or ',')

101787 A # : *reference* comment indicates the location(s) of the **msgid** string in the source files, in
101788 *pathname1:linenumber1 [pathname2:linenumber2 ...]*

101789 format. They can be added, as might "# ." prefixed additional comments of unspecified format,
101790 by the *xgettext* utility. All comments that do not begin with "# ," are informative only and shall
101791 be silently ignored by the *msgfmt* utility. In "# ," comments the following values for *flag* can be

101792 specified:

101793 **fuzzy** This flag indicates that the **msgid** string might not be a correct translation at this
 101794 point in time. Only the translator can judge if the translation requires further
 101795 modification or is acceptable as is. Once satisfied with the translation, the
 101796 translator should remove this **fuzzy** flag. If this flag is specified, the *msgfmt* utility
 101797 shall not generate the entry for the next following **msgid** in the output message
 101798 catalog, unless the **-f** option is specified. If other flag comments are specified
 101799 between **fuzzy** and the **msgid**, the behavior is unspecified.

101800 **c-format**

101801 **no-c-format** The **c-format** flag indicates that the next following **msgid** string contains a *printf()*
 101802 format string. When the **c-format** flag is given and the **-c** and **-v** options are
 101803 specified, the *msgfmt* utility shall perform additional tests to check the validity of
 101804 the translation (see OPTIONS); these additional tests may also be performed if
 101805 neither **c-format** nor **no-c-format** is given. When the **no-c-format** flag is given for a
 101806 string, no additional checks shall be performed for the string. When both the **c-**
 101807 **format** and the **no-c-format** flags are given, the last flag specified takes precedence.

101808 EXIT STATUS

101809 The following exit values shall be returned:

101810 0 Successful completion.

101811 >0 An error occurred.

101812 CONSEQUENCES OF ERRORS

101813 The *msgfmt* utility need not continue processing later *pathname* operands when an error
 101814 condition that affects the exit status is detected. It is unspecified whether a messages object file is
 101815 written when checks performed for the **-c** and **-v** options fail.

101816 APPLICATION USAGE

101817 The *xgettext* utility can be used to create template dot-po files from C-language source files.

101818 Installing messages object files for the POSIX or C locale is not recommended, since they may be
 101819 ignored for the sake of efficiency.

101820 The first section for each domain in a dot-po file should include a header containing a

101821 "charset=*codeset*"

101822 specification. If this specification is omitted, message conversions in the *gettext* family of
 101823 functions and in the *gettext* and *ngettext* utilities may fail.

101824 The **msgid_plural** directive's *untranslated_string_plural* string comes from the *msgid_plural*
 101825 arguments in calls to the *ngettext()*, *ngettext_l()*, *dngettext()*, *dngettext_l()*, *dcngettext()*, and
 101826 *dcngettext_l()* functions when a prototype dot-po file is created by the *xgettext* utility. These
 101827 strings (and the *msgid_plural* operands in calls to the *ngettext* utility) can provide context when a
 101828 translator is modifying a template dot-po file into a dot-po file for a specific language. These
 101829 functions and the *ngettext* utility do not try to match the *msgid_plural* arguments or operands
 101830 with anything in a messages object file; they only match the *msgid* arguments and operands.

101831 Unlike shell command language strings, double-quoted strings in dot-po files cannot contain a
 101832 literal <newline> character.

101833 **EXAMPLES**

```
101834     In this example, module1.po and module2.po are portable messages object source files.
101835     $ cat module1.po
101836     # default domain "messages"
101837     msgid ""
101838     msgstr "charset=utf-8"
101839     msgid "msg 1"
101840     msgstr "msg 1 translation"
101841     #
101842     domain "help_domain"
101843     msgid ""
101844     msgstr "charset=utf-8"
101845     msgid "help 2"
101846     msgstr "help 2 translation"
101847     #
101848     domain "error_domain"
101849     msgid ""
101850     msgstr "charset=utf-8"
101851     msgid "error 3"
101852     msgstr "error 3 translation"
101853
101853     $ cat module2.po
101854     # default domain "messages"
101855     msgid ""
101856     msgstr "charset=utf-8"
101857     msgid "mesg 4"
101858     msgstr "mesg 4 translation"
101859     #
101860     domain "error_domain"
101861     msgid ""
101862     msgstr "charset=utf-8"
101863     #, c-format
101864     msgid "error 5 %s"
101865     msgstr "error 5 translation %s"
101866     #
101867     domain "window_domain"
101868     msgid ""
101869     msgstr "charset=utf-8"
101870     msgid "window 6"
101871     msgstr "window 6 translation"
101872
101872     $ cat module3.po
101873     # default domain "messages"
101874     # header will be used for the whole output file in the third example
101875     msgid ""
101876     msgstr "charset=utf-8"
101877     msgid "info 0"
101878     msgstr "info 0 translation"
101879
101879     $ cat opt_debug.po
101880     #
101881     domain "debug_domain"
101882     msgid "debug 8"
```

101883 **msgstr "debug 8 translation"**

101884 The following command will produce the output files **messages.mo**, **help_domain.mo**, and
101885 **error_domain.mo**:

101886 \$ msgfmt -S module1.po

101887 The following command will produce the output files **messages.mo**, **help_domain.mo**,
101888 **error_domain.mo**, and **window_domain.mo**:

101889 \$ msgfmt -S module1.po module2.po

101890 The following command will produce the output file **hello.mo**:

101891 \$ msgfmt -o hello.mo module3.po opt_debug.po

101892 **RATIONALE**

101893 Some implementations are less strict about the format of dot-po files and simply treat all
101894 occurrences of one or more white space characters as a separator. The format described in this
101895 standard is accepted by all known implementations.

101896 In some implementations, duplicate **msgid** directives within a domain are ignored, and only an
101897 entry for the first **msgid** directive and the following **msgid**, **msgid_plural**, **msgstr**, or
101898 **msgstr[index]** directives is created. However, some implementations consider duplicate **msgid**
101899 directives within a domain to be an error and do not produce output at all. Consequently this
101900 standard does not specify the behavior of *msgfmt* if duplicate **msgid** directives are encountered
101901 within one domain.

101902 **FUTURE DIRECTIONS**

101903 None.

101904 **SEE ALSO**

101905 *gettext*, *xgettext*

101906 XSH *fprintf()*, *gettext*

101907 **CHANGE HISTORY**

101908 First released in Issue 8.

101909

102281 **NAME**

102282 ngettext — retrieve text string from messages object

102283 **SYNOPSIS**

102284 ngettext [-e|-E] [-d *textdomain*] [*textdomain*] *msgid msgid_plural n*

102285 **DESCRIPTION**

102286 Refer to *gettext*.

102772 *LC_CTYPE* Determine the locale for the interpretation of sequences of bytes of text data as
 102773 characters (for example, single-byte as opposed to multi-byte characters in
 102774 arguments).

102775 *LC_MESSAGES*
 102776 Determine the locale that should be used to affect the format and contents of
 102777 diagnostic messages written to standard error.

102778 XSI *NLSPATH* Determine the location of message catalogs for the processing of *LC_MESSAGES*.

102779 *PATH* Determine the search path that is used to locate the utility to be invoked. See XBD
 102780 [Chapter 8](#) (on page 155).

102781 **ASYNCHRONOUS EVENTS**

102782 The *nohup* utility shall take the standard action for all signals except that SIGHUP shall be
 102783 ignored.

102784 **STDOUT**

102785 If the standard output is not a terminal, the standard output of *nohup* shall be the standard
 102786 output generated by the execution of the *utility* specified by the operands. Otherwise, nothing
 102787 shall be written to the standard output.

102788 **STDERR**

102789 If the standard output is a terminal, a message shall be written to the standard error, indicating
 102790 the name of the file to which the output is being appended. The name of the file shall be either
 102791 **nohup.out** or **\$HOME/nohup.out**.

102792 **OUTPUT FILES**

102793 Output written by the named utility is appended to the file **nohup.out** (or **\$HOME/nohup.out**),
 102794 if the conditions hold as described in the DESCRIPTION.

102795 **EXTENDED DESCRIPTION**

102796 None.

102797 **EXIT STATUS**

102798 The following exit values shall be returned:

102799 126 The utility specified by *utility* was found but could not be invoked.

102800 127 An error occurred in the *nohup* utility or the utility specified by *utility* could not be
 102801 found.

102802 Otherwise, the exit status of *nohup* shall be that of the utility specified by the *utility* operand.

102803 **CONSEQUENCES OF ERRORS**

102804 Default.

102805 **APPLICATION USAGE**

102806 The *command*, *env*, *nice*, *nohup*, *time*, *timeout*, and *xargs* utilities have been specified to use exit +
 102807 code 127 if an error occurs so that applications can distinguish “failure to find a utility” from
 102808 “invoked utility exited with an error indication”. The value 127 was chosen because it is not
 102809 commonly used for other meanings; most utilities use small values for “normal error
 102810 conditions” and the values above 128 can be confused with termination due to receipt of a
 102811 signal. The value 126 was chosen in a similar manner to indicate that the utility could be found,
 102812 but not invoked. Some scripts produce meaningful error messages differentiating the 126 and
 102813 127 cases. The distinction between exit codes 126 and 127 is based on KornShell practice that
 102814 uses 127 when all attempts to *exec* the utility fail with [ENOENT], and uses 126 when any
 102815 attempt to *exec* the utility fails for any other reason.

106666 **NAME**

106667 readlink — display the contents of a symbolic link

106668 **SYNOPSIS**106669 readlink [-n] *file*106670 **DESCRIPTION**

106671 If the *file* operand names a symbolic link, the *readlink* utility shall not follow the symbolic link
 106672 when resolving *file* and shall write the contents of the symbolic link to standard output. If the **-n**
 106673 option is not specified, the output to standard output shall be followed by a <newline>
 106674 character.

106675 If *file* does not name a symbolic link, *readlink* shall write a diagnostic message to standard error
 106676 and exit with non-zero status.

106677 **OPTIONS**106678 The *readlink* utility shall conform to XBD [Section 12.2](#) (on page 201).

106679 The following option shall be supported:

106680 **-n** Do not output a trailing <newline> character.106681 **OPERANDS**

106682 The following operand shall be supported:

106683 *file* A pathname of a symbolic link to be read.106684 **STDIN**

106685 Not used.

106686 **INPUT FILES**

106687 None.

106688 **ENVIRONMENT VARIABLES**106689 The following environment variables shall affect the execution of *readlink*:

106690 **LANG** Provide a default value for the internationalization variables that are unset or null.
 106691 (See XBD [Section 8.2](#) (on page 157) for the precedence of internationalization
 106692 variables used to determine the values of locale categories.)

106693 **LC_ALL** If set to a non-empty string value, override the values of all the other
 106694 internationalization variables.

106695 **LC_CTYPE** Determine the locale for the interpretation of sequences of bytes of text data as
 106696 characters (for example, single-byte as opposed to multi-byte characters in
 106697 arguments and input files).

106698 **LC_MESSAGES**

106699 Determine the locale that should be used to affect the format and contents of
 106700 diagnostic messages written to standard error.

106701 XSI **NLSPATH** Determine the location of messages objects and message catalogs.106702 **ASYNCHRONOUS EVENTS**

106703 Default.

106704 **STDOUT**

106705 See DESCRIPTION.

106706 **STDERR**

106707 The standard error shall be used only for diagnostic messages.

106708 **OUTPUT FILES**

106709 None.

106710 **EXTENDED DESCRIPTION**

106711 None.

106712 **EXIT STATUS**

106713 The following exit values shall be returned:

106714 0 Successful completion.

106715 >0 An error occurred.

106716 **CONSEQUENCES OF ERRORS**

106717 Default.

106718 **APPLICATION USAGE**

106719 None.

106720 **EXAMPLES**

106721 None.

106722 **RATIONALE**

106723 The *readlink* utility was added because using *ls -l* to obtain the contents of a symbolic link is difficult if the output includes more than one occurrence of the string " -> ".

106725 The *-f* option found in many implementations was not included, as the *realpath* utility provides equivalent functionality with a choice of behaviors.

106727 **FUTURE DIRECTIONS**

106728 None.

106729 **SEE ALSO**

106730 *ln*, *ls*, *realpath*

106731 XBD [Chapter 8](#) (on page 155), [Section 12.2](#) (on page 201)

106732 XSH [readlink\(\)](#)

106733 **CHANGE HISTORY**

106734 First released in Issue 8.

106735

NAME

106736 realpath — resolve a pathname
106737

SYNOPSIS

106738 realpath [-E|-e] *file*
106739

DESCRIPTION

106740 The *realpath* utility shall canonicalize the pathname specified by the *file* operand as follows:
106741

106742 If a call to the *realpath*() function with the specified pathname as its first argument would
106743 succeed, the canonicalized pathname shall be the pathname that would be returned by that
106744 *realpath*() call. Otherwise:

- 106745 • If the **-e** option is specified, the canonicalization shall fail.
- 106746 • If the **-E** option is specified, then if a call to the *realpath*() function with the specified
106747 pathname as its first argument would encounter an error condition other than [ENOENT],
106748 the canonicalization shall fail; if the call would encounter an [ENOENT] error, *realpath* shall
106749 expand all symbolic links that would be encountered in an attempt to resolve the specified
106750 pathname using the algorithm specified in XBD [Section 4.14](#) (on page 93), except that any
106751 trailing <slash> characters that are not also leading <slash> characters shall be ignored. If
106752 this expansion succeeds and the path prefix of the expanded pathname resolves to an
106753 existing directory, the canonicalized pathname shall be the expanded pathname. In all
106754 other cases, the canonicalization shall fail. If the expanded pathname is not empty, does not
106755 begin with a <slash>, and has exactly one pathname component, it shall be treated as if it
106756 had a path prefix of ". /".
- 106757 • If no options are specified, *realpath* shall canonicalize the specified pathname in an
106758 unspecified manner such that the resulting absolute pathname does not contain any
106759 components that refer to files of type symbolic link and does not contain any components
106760 that are dot or dot-dot.

106761 Upon successful canonicalization, *realpath* shall write the canonicalized pathname, followed by a
106762 <newline> character, to standard output.

106763 If canonicalization fails, or the canonicalized pathname is empty, nothing shall be written to
106764 standard output, a diagnostic message shall be written to standard error, and *realpath* shall exit
106765 with non-zero status.

OPTIONS

106766 The *realpath* utility shall conform to XBD [Section 12.2](#) (on page 201).
106767

106768 The following options shall be supported:

- 106769 **-E** Do not treat it as an error if attempting to resolve the last component of the
106770 canonicalized form of the *file* operand results in an [ENOENT] error condition.
- 106771 **-e** Treat it as an error if attempting to resolve the last component of the canonicalized
106772 form of the *file* operand results in an [ENOENT] error condition.

106773 Specifying more than one of the mutually-exclusive options **-E** and **-e** shall not be considered an
106774 error. The last option specified shall determine the behavior of the utility.

OPERANDS

106775 The following operand shall be supported:
106776

106777 *file* A pathname to be canonicalized.

106778 **STDIN**
106779 Not used.

106780 **INPUT FILES**
106781 None.

106782 **ENVIRONMENT VARIABLES**
106783 The following environment variables shall affect the execution of *realpath*:

106784 *LANG* Provide a default value for the internationalization variables that are unset or null.
106785 (See XBD [Section 8.2](#) (on page 157) for the precedence of internationalization
106786 variables used to determine the values of locale categories.)

106787 *LC_ALL* If set to a non-empty string value, override the values of all the other
106788 internationalization variables.

106789 *LC_CTYPE* Determine the locale for the interpretation of sequences of bytes of text data as
106790 characters (for example, single-byte as opposed to multi-byte characters in
106791 arguments and input files).

106792 *LC_MESSAGES*
106793 Determine the locale that should be used to affect the format and contents of
106794 diagnostic messages written to standard error.

106795 XSI *NLSPATH* Determine the location of messages objects and message catalogs.

106796 **ASYNCHRONOUS EVENTS**
106797 Default.

106798 **STDOUT**
106799 See DESCRIPTION.

106800 **STDERR**
106801 The standard error shall be used only for diagnostic messages.

106802 **OUTPUT FILES**
106803 None.

106804 **EXTENDED DESCRIPTION**
106805 None.

106806 **EXIT STATUS**
106807 The following exit values shall be returned:

106808 0 Successful completion.
106809 >0 An error occurred.

106810 **CONSEQUENCES OF ERRORS**
106811 Default.

106812 **APPLICATION USAGE**

106813 If neither the `-e` nor the `-E` option is specified, some implementations behave as if `-e` had been
 106814 specified and others as if `-E` had been specified, but there are also implementations where the
 106815 behavior differs from both of these. For example, the *mksh* shell has an internal implementation
 106816 of *realpath* that canonicalizes `/dir/regular_file/..` to `/dir`, whereas the *realpath()* function would
 106817 return an [ENOTDIR] error in this case. Portable applications should always specify either `-e` or
 106818 `-E`.

106819 **EXAMPLES**

106820 None.

106821 **RATIONALE**

106822 The *realpath* utility was added in preference to a `-f` option found in some implementations of the
 106823 *readlink* utility because it allows the application to specify whether or not a missing final
 106824 component is to be treated as an error.

106825 The behavior with the `-E` option when *file* does not resolve (with symbolic links followed) to an
 106826 existing file is not the same as simply calling *realpath()* with the path prefix of the *file* operand
 106827 and writing the resulting pathname, a `<slash>`, and the last component of *file* to standard output.
 106828 For example, if `/tmp/nofile` does not exist, and *file* is `A/B` where `A` is an existing directory and `B`
 106829 is a symbolic link to `/tmp/nofile`, *realpath* with `-E` will output `/tmp/nofile`, but if `B` is a symbolic
 106830 link to `/tmp/nofile/foo`, *realpath* with `-E` will treat this as an error. In both cases
 106831 *realpath("A/B")* would fail with *errno* set to [ENOENT]. Even though *realpath("A")*
 106832 would succeed, in neither case is anything ending `/B` the result.

106833 Trailing `<slash>` characters (that follow a non-`<slash>`) are handled differently with `-E` than with
 106834 `-e`. With `-e` they are handled as for the *realpath()* function. With `-E` they are sometimes
 106835 effectively ignored, and they are never included in the output. For example, if `/tmp/nofile` does
 106836 not exist and `/tmp/regfile` is an existing regular file:

```
106837 $ realpath -E /tmp/nofile/  
106838 /tmp/nofile  
106839 $ realpath -E /tmp/regfile/  
106840 realpath: /tmp/regfile/: Not a directory
```

106841 Although the behavior of the *realpath* utility is specified by reference to the *realpath()* function,
 106842 which is part of the XSI option, non-XSI implementations that do not support *realpath()* are
 106843 nevertheless required to implement *realpath* in accordance with the requirements described in
 106844 this standard for *realpath()*.

106845 **FUTURE DIRECTIONS**

106846 None.

106847 **SEE ALSO**

106848 *ln*, *ls*, *pwd*, *readlink*

106849 XBD Chapter 8 (on page 155), Section 12.2 (on page 201)

106850 XSH Section 2.3 (on page 475), *realpath()*

106851 **CHANGE HISTORY**

106852 First released in Issue 8.

106853

110894 127 The utility specified by *utility* could not be found.

110895 CONSEQUENCES OF ERRORS

110896 Default.

110897 APPLICATION USAGE

110898 The *command*, *env*, *nice*, *nohup*, *time*, *timeout*, and *xargs* utilities have been specified to use exit +
 110899 code 127 if an error occurs so that applications can distinguish “failure to find a utility” from
 110900 “invoked utility exited with an error indication”. The value 127 was chosen because it is not
 110901 commonly used for other meanings; most utilities use small values for “normal error
 110902 conditions” and the values above 128 can be confused with termination due to receipt of a
 110903 signal. The value 126 was chosen in a similar manner to indicate that the utility could be found,
 110904 but not invoked. Some scripts produce meaningful error messages differentiating the 126 and
 110905 127 cases. The distinction between exit codes 126 and 127 is based on KornShell practice that
 110906 uses 127 when all attempts to *exec* the utility fail with [ENOENT], and uses 126 when any
 110907 attempt to *exec* the utility fails for any other reason.

110908 EXAMPLES

110909 It is frequently desirable to apply *time* to pipelines or lists of commands. This can be done by
 110910 placing pipelines and command lists in a single file; this file can then be invoked as a utility, and
 110911 the *time* applies to everything in the file.

110912 Alternatively, the following command can be used to apply *time* to a complex command:

```
110913 time sh -c 'complex-command-line'
```

110914 RATIONALE

110915 When the *time* utility was originally proposed to be included in the ISO POSIX-2:1993 standard,
 110916 questions were raised about its suitability for inclusion on the grounds that it was not useful for
 110917 conforming applications, specifically:

- 110918 • The underlying CPU definitions from the System Interfaces volume of POSIX.1-202x are
 110919 vague, so the numeric output could not be compared accurately between systems or even
 110920 between invocations.
- 110921 • The creation of portable benchmark programs was outside the scope this volume of
 110922 POSIX.1-202x.

110923 However, *time* does fit in the scope of user portability. Human judgement can be applied to the
 110924 analysis of the output, and it could be very useful in hands-on debugging of applications or in
 110925 providing subjective measures of system performance. Hence it has been included in this
 110926 volume of POSIX.1-202x.

110927 The default output format has been left unspecified because historical implementations differ
 110928 greatly in their style of depicting this numeric output. The `-p` option was invented to provide
 110929 scripts with a common means of obtaining this information.

110930 In the KornShell, *time* is a shell reserved word that can be used to time an entire pipeline, rather
 110931 than just a simple command. The POSIX definition has been worded to allow this
 110932 implementation. Consideration was given to invalidating this approach because of the historical
 110933 model from the C shell and System V shell. However, since the System V *time* utility historically
 110934 has not produced accurate results in pipeline timing (because the constituent processes are not
 110935 all owned by the same parent process, as allowed by POSIX), it did not seem worthwhile to
 110936 break historical KornShell usage.

110937 The term *utility* is used, rather than *command*, to highlight the fact that shell compound
 110938 commands, pipelines, special built-ins, and so on, cannot be used directly. However, *utility*
 110939 includes user application programs and shell scripts, not just the standard utilities.

110958 **NAME**

110959 timeout — execute a utility with a time limit

110960 **SYNOPSIS**110961 timeout [-fp] [-k *time*] [-s *signal_name*] *duration utility [argument...]*110962 **DESCRIPTION**

110963 The *timeout* utility shall execute the utility named by the *utility* operand, with arguments
 110964 supplied as the *argument* operands (if any), in a child process. If the value of the *duration*
 110965 operand is non-zero and the child process has not terminated after the specified time period,
 110966 *timeout* shall send the signal specified by the *-s* option, or the SIGTERM signal if *-s* is not given.

110967 If the *-f* option is specified, the signal shall be sent only to the child process. Otherwise, it is
 110968 implementation defined which one of the following methods is used to signal additional
 110969 processes:

- 110970 • The *timeout* utility ensures it is a process group leader before creating the child process
 110971 which executes the utility, in which case it shall send the signal to its process group.
- 110972 • The *timeout* utility arranges for any descendents of the child process that are orphaned to
 110973 have their parent process changed to the *timeout* utility, in which case the signal shall be
 110974 sent to the child process and all of its descendents.

110975 If the subsequent wait status of the child process shows that it was stopped by a signal, a
 110976 SIGCONT signal shall also be sent in the same manner as the first signal; otherwise, a SIGCONT
 110977 signal may be sent in the same manner.

110978 If the *-k* option is specified, and the child process created to execute the utility still has not
 110979 terminated after the time period specified by the *time* option-argument has elapsed since the first
 110980 signal was sent, *timeout* shall send a SIGKILL signal in the same manner as the first signal. If
 110981 *timeout* receives a signal and propagates it to the child process (see ASYNCHRONOUS EVENTS
 110982 below), this shall be treated as the first signal.

110983 **OPTIONS**110984 The *timeout* utility shall conform to XBD [Section 12.2](#) (on page 201).

110985 The following options shall be supported:

- 110986 **-f** Only time out the utility itself, not its descendents.
- 110987 **-k *time*** Send a SIGKILL signal if the child process created to execute the utility has not
 110988 terminated after the time period specified by *time* has elapsed since the first signal
 110989 was sent. The value of *time* shall be interpreted as specified for the *duration*
 110990 operand (see OPERANDS below).
- 110991 **-p** Always preserve (mimic) the wait status of the executed utility, even if the time
 110992 limit was reached.
- 110993 **-s *signal_name***
 110994 Specify the signal to send when the time limit is reached, using one of the symbolic
 110995 names defined in the **<signal.h>** header. Values of *signal_name* shall be recognized
 110996 in a case-independent fashion, without the SIG prefix. By default, SIGTERM shall
 110997 be sent.

110998 **OPERANDS**

110999 The following operands shall be supported:

- 111000 *duration* The maximum amount of time to allow the utility to run, specified as a decimal
 111001 number with an optional decimal fraction and an optional suffix, which can be:

111002		s seconds
111003		m minutes
111004		h hours
111005		d days
111006		If a decimal fraction is present, the application shall ensure that it is separated from the units by a <period>. If no suffix is present, the value shall specify seconds.
111007		
111008		If the value is zero, <i>timeout</i> shall not enforce a time limit.
111009	<i>utility</i>	The name of a utility that is to be executed. If the <i>utility</i> operand names any of the special built-in utilities in Section 2.14 (on page 2382), the results are undefined.
111010		
111011	<i>argument</i>	Any string to be supplied as an argument when executing the utility named by the <i>utility</i> operand.
111012		
111013	STDIN	
111014		Not used.
111015	INPUT FILES	
111016		None.
111017	ENVIRONMENT VARIABLES	
111018		The following environment variables shall affect the execution of <i>timeout</i> :
111019	<i>LANG</i>	Provide a default value for the internationalization variables that are unset or null. (See XBD Section 8.2 (on page 157) for the precedence of internationalization variables used to determine the values of locale categories.)
111020		
111021		
111022	<i>LC_ALL</i>	If set to a non-empty string value, override the values of all the other internationalization variables.
111023		
111024	<i>LC_CTYPE</i>	Determine the locale for the interpretation of sequences of bytes of text data as characters (for example, single-byte as opposed to multi-byte characters in arguments and input files).
111025		
111026		
111027	<i>LC_MESSAGES</i>	
111028		Determine the locale that should be used to affect the format and contents of diagnostic messages written to standard error.
111029		
111030	XSI <i>NLSPATH</i>	Determine the location of messages objects and message catalogs.
111031	<i>PATH</i>	Determine the search path that is used to locate the utility to be executed. See XBD Section 8.3 (on page 162).
111032		
111033	ASYNCHRONOUS EVENTS	
111034		The default behavior specified in Section 1.4 (on page 2328) shall apply, except that:
111035		• The <i>timeout</i> utility shall ignore SIGTTIN and SIGTTOU signals.
111036		• The <i>timeout</i> utility may alter the disposition of SIGALRM if the inherited disposition was for it to be ignored.
111037		
111038		• If the signal specified with the <i>-s</i> option, or any signal whose default action is to terminate the process, is delivered to the <i>timeout</i> utility, then unless the signal is SIGKILL or SIGSTOP, the <i>timeout</i> utility shall immediately send the same signal to the process or processes to which it would send a signal when the time limit is reached. If the delivered signal is SIGALRM, <i>timeout</i> may behave as if the time limit had been reached instead of sending SIGALRM.
111039		
111040		
111041		
111042		
111043		

111044 • If the **-f** option is not specified, then if *timeout* sends a signal to its process group, it shall
111045 briefly change the disposition of that signal to ignored while it sends the signal, so that it
111046 does not receive the signal itself.

111047 With the single exception of the signal specified with the **-s** option, or SIGTERM if **-s** is not
111048 used, all signal dispositions inherited by the utility specified by the *utility* operand shall be the
111049 same as the disposition that *timeout* inherited.

111050 **STDOUT**

111051 Not used.

111052 **STDERR**

111053 The standard error shall be used only for diagnostic messages.

111054 **OUTPUT FILES**

111055 None.

111056 **EXTENDED DESCRIPTION**

111057 None.

111058 **EXIT STATUS**

111059 If the **-p** option is not specified and the time limit was reached:

111060 • If the **-k** option was not specified or the utility terminated before the time period specified
111061 by the *time* option-argument elapsed since the first signal was sent, the exit status shall be
111062 124.

111063 • If the **-k** option was specified and the SIGKILL signal was sent, it is unspecified whether
111064 the exit status is 124 or the behavior is as if the **-p** option was specified.

111065 Otherwise, if the executed utility terminated by exiting, the exit status of *timeout* shall be that of
111066 the utility; if the utility was terminated by a signal, *timeout* shall terminate itself with the same
111067 signal while ensuring that a core image is not created.

111068 If an error occurs, the following exit values shall be returned:

111069 125 An error other than the two described below occurred.

111070 126 The utility specified by *utility* was found but could not be executed.

111071 127 The utility specified by *utility* could not be found.

111072 **CONSEQUENCES OF ERRORS**

111073 Default.

111074 **APPLICATION USAGE**

111075 Unlike the *kill* utility, the **-s** option of *timeout* is not required to accept the symbolic name 0 to
111076 represent signal value zero.

111077 When the value of *duration* is zero, *timeout* does not time out the utility, but it does still perform
111078 signal propagation (including to descendants of the utility if **-f** is not specified).

111079 Regardless of locale, the `<period>` character (the decimal-point character of the POSIX locale) is
111080 the decimal-point character recognized in the *duration* operand and the *time* option-argument.

111081 The *command*, *env*, *nice*, *nohup*, *time*, *timeout*, and *xargs* utilities have been specified to use exit
111082 code 127 if a utility to be invoked cannot be found, so that applications can distinguish “failure
111083 to find a utility” from “invoked utility exited with an error indication”. The value 127 was
111084 chosen because it is not commonly used for other meanings; most utilities use small values for
111085 “normal error conditions” and the values above 128 can be confused with termination due to
111086 receipt of a signal. The value 126 was chosen in a similar manner to indicate that the utility could

111087 be found, but not invoked. Some scripts produce meaningful error messages differentiating the
111088 126 and 127 cases. The distinction between exit codes 126 and 127 is based on KornShell practice
111089 that uses 127 when all attempts to *exec* the utility fail with [ENOENT], and uses 126 when any
111090 attempt to *exec* the utility fails for any other reason. The *timeout* utility extends these special exit
111091 codes to 125 and 124, with the meanings described in EXIT STATUS. A *timeout* exit status below
111092 124 can only result from passing through the exit status of the executed utility.

111093 EXAMPLES

111094 None.

111095 RATIONALE

111096 Some *timeout* implementations make themselves a process group leader (when *-f* is not used) in
111097 order to be able to send signals to descendants of the child process. However, using this method
111098 means that any descendants which change their process group do not receive the signal. To
111099 ensure all descendants receive the signal, some implementations instead make use of a feature
111100 whereby descendants that are orphaned have their parent process changed to the *timeout*
111101 utility—that is, *timeout* becomes their “reaper”—together with the ability of a reaper to send a
111102 signal to all of its descendants.

111103 Some historical *timeout* implementations exited with status *128+signal_number* when the child
111104 process was terminated by a signal before the time limit was reached (or when *-p* was used).
111105 This is reasonable when *timeout* is invoked from a shell which sets *\$?* to *128+signal_number*, but
111106 not all shells do that. In particular, the KornShell sets *\$?* to *256+signal_number* and so an exit
111107 status of *128+signal_number* from *timeout* would be misleading. In order to avoid any possible
111108 ambiguity, this standard requires that *timeout* mimics the wait status of the child process by
111109 terminating itself with the same signal. When it does this it needs to ensure that it does not
111110 create a core image, otherwise it could overwrite one created by the invoked utility.

111111 The *timeout* utility ignores SIGTTIN and SIGTTOU so that if the utility it executes reads from or
111112 writes to the controlling terminal and this generates a SIGTTIN or SIGTTOU for the process
111113 group, *timeout* will not be stopped by the signal and can still time out the utility.

111114 Some historical *timeout* implementations always set the disposition for SIGTTIN and SIGTTOU
111115 in the child process to default, even if these signals were inherited as ignored. This could result
111116 in processes being stopped unexpectedly. Likewise, they did not ensure that for signals they
111117 caught, the disposition inherited by the executed utility was the same as the disposition that was
111118 inherited by *timeout*. This meant that, for example, if *timeout* was used in a script that was run
111119 with *nohup*, the utility executed by *timeout* would unexpectedly not be protected from SIGHUP.
111120 This standard requires that all signal dispositions inherited by the utility specified by the *utility*
111121 operand are the same as the disposition that *timeout* inherited, with the single exception of the
111122 signal that *timeout* sends when the time limit is reached, which needs to be inherited as default
111123 in order for the timeout to take effect (without resorting to SIGKILL if *-k* is specified).

111124 Some historical *timeout* implementations only propagated a subset of the signals whose default
111125 action is to terminate the process to the child process if one was delivered to the *timeout* utility.
111126 Propagating these signals is beneficial, as otherwise termination of the *timeout* utility by a signal
111127 results in the utility it executed being left running indefinitely (unless it also received the signal,
111128 for example a terminal-generated SIGINT). There is no reason to select a subset of these signals
111129 to be propagated, therefore this standard requires them all to be propagated (except SIGKILL,
111130 which cannot). In the event that a user wants to prevent the utility being timed out, sending
111131 *timeout* a SIGKILL can be used for this purpose.

111132 **FUTURE DIRECTIONS**

111133 None.

111134 **SEE ALSO**111135 *kill*111136 XBD [Chapter 8](#) (on page 155), [Section 12.2](#) (on page 201), [<signal.h>](#) (on page 329)111137 **CHANGE HISTORY**

111138 First released in Issue 8.

111139

116531 **EXTENDED DESCRIPTION**

116532 None.

116533 **EXIT STATUS**

116534 The following exit values shall be returned:

- 116535 0 All invocations of *utility* returned exit status zero.
- 116536 1-125 A command line meeting the specified requirements could not be assembled, one or more of the invocations of *utility* returned a non-zero exit status, or some other error occurred.
- 116537
- 116538
- 116539 126 The utility specified by *utility* was found but could not be invoked.
- 116540 127 The utility specified by *utility* could not be found.

116541 **CONSEQUENCES OF ERRORS**

116542 If a command line meeting the specified requirements cannot be assembled, the utility cannot be invoked, an invocation of the utility is terminated by a signal, or an invocation of the utility exits with exit status 255, the *xargs* utility shall write a diagnostic message and exit without processing any remaining input.

116543

116544

116545

116546 **APPLICATION USAGE**

116547 The 255 exit status allows a utility being used by *xargs* to tell *xargs* to terminate if it knows no further invocations using the current data stream will succeed. Thus, *utility* should explicitly *exit* with an appropriate value to avoid accidentally returning with 255.

116548

116549

116550 Note that since input is parsed as lines, <blank> characters separate arguments, and <backslash>, <apostrophe>, and double-quote characters are used for quoting, if *xargs* is used to bundle the output of commands like *find dir -print* or *ls* into commands to be executed, unexpected results are likely if any filenames contain <blank>, <newline>, or quoting characters. This can be solved by using *find* to call a script that converts each file found into a quoted string that is then piped to *xargs*, but in most cases it is preferable just to have *find* do the argument aggregation itself by using *-exec* with a '+' terminator instead of ';' . Note that the quoting rules used by *xargs* are not the same as in the shell. They were not made consistent here because existing applications depend on the current rules. An easy (but inefficient) method that can be used to transform input consisting of one argument per line into a quoted form that *xargs* interprets correctly is to precede each non-<newline> character with a <backslash>. More efficient alternatives are shown in Example 2 and Example 5 below.

116551

116552

116553

116554

116555

116556

116557

116558

116559

116560

116561

116562 On implementations with a large value for {ARG_MAX}, *xargs* may produce command lines longer than {LINE_MAX}. For invocation of utilities, this is not a problem. If *xargs* is being used to create a text file, users should explicitly set the maximum command line length with the *-s* option.

116563

116564

116565

116566 The *command*, *env*, *nice*, *nohup*, *time*, *timeout*, and *xargs* utilities have been specified to use exit + code 127 if an error occurs so that applications can distinguish "failure to find a utility" from "invoked utility exited with an error indication". The value 127 was chosen because it is not commonly used for other meanings; most utilities use small values for "normal error conditions" and the values above 128 can be confused with termination due to receipt of a signal. The value 126 was chosen in a similar manner to indicate that the utility could be found, but not invoked. Some scripts produce meaningful error messages differentiating the 126 and 127 cases. The distinction between exit codes 126 and 127 is based on KornShell practice that uses 127 when all attempts to *exec* the utility fail with [ENOENT], and uses 126 when any attempt to *exec* the utility fails for any other reason.

116567

116568

116569

116570

116571

116572

116573

116574

116575

116690 **NAME**116691 `xgettext` — extract `gettext` call strings from C-language source files (DEVELOPMENT)116692 **SYNOPSIS**

```
116693 CD xgettext [-j] [-n] [-d default-domain] [-K keyword-spec]...
116694 [-p pathname] file...
116695 xgettext -a [-n] [-d default-domain] [-p pathname]
116696 [-x exclude-file] file...
116697
```

116698 **DESCRIPTION**

116699 The `xgettext` utility shall automate the creation of portable messages object source files (dot-po
 116700 files). A dot-po file shall contain copies of string literals that are found in C-language source
 116701 code in files specified by *file* operands. The dot-po file can be used as input to the `msgfmt` utility,
 116702 to produce a messages object file that can be used by applications.

116703 The `xgettext` utility shall write `msgid` argument strings that are passed as string literals in
 116704 `gettext()`, `gettext_l()`, `ngettext()`, and `ngettext_l()` calls in C-language source code to the default
 116705 output file; this file shall be named **messages.po** unless it is changed by the `-d` option. The
 116706 `xgettext` utility shall also write `msgid` argument strings that are passed as string literals in
 116707 `dcgettext()`, `dcgettext_l()`, `dcngettext()`, `dcngettext_l()`, `dgettext()`, `dgettext_l()`, `dngettext()`, and
 116708 `dngettext_l()` calls either to the default output file or to the output file `domainname.po` where
 116709 `domainname` is the first parameter to the call; it is implementation-defined which of those output
 116710 files is used. A **msgid** directive shall precede each `msgid` argument string. For the functions that
 116711 have a `msgid_plural` argument, a **msgid_plural** directive followed by that argument string shall
 116712 also be written directly after the corresponding **msgid** directive. A **msgstr** directive or
 116713 **msgstr[index]** directives with an empty string shall be written after the corresponding **msgid** or
 116714 **msgid_plural** directive, respectively. The function names that `xgettext` searches for can be
 116715 changed using the `-K` option.

116716 The first directive in each created dot-po file shall be a **domain** directive giving the associated
 116717 domain name, except that this directive is optional in the default output file.

116718 If the `-p pathname` option is specified, `xgettext` shall create the dot-po files in the `pathname`
 116719 directory. Otherwise, the dot-po files shall be created in the current working directory.

116720 The **msgid** values shall be in the same order that the strings are extracted from each *file* and
 116721 subsections with duplicate **msgid** values shall be written to the dot-po files as comment lines.

116722 **OPTIONS**

116723 The `xgettext` utility shall conform to XBD [Section 12.2](#) (on page 201).

116724 The following options shall be supported:

116725 **-a** Extract all strings, not just those found in calls to `gettext` family functions. Only one
 116726 dot-po file shall be created.

116727 **-d default-domain**
 116728 Name the default output file `default-domain.po` instead of **messages.po**.

116729 **-j** Join messages from C-language source files with existing dot-po files. For each
 116730 dot-po file that `xgettext` writes messages to, if the file does not exist, it shall be
 116731 created. New messages shall be appended but any subsections with duplicate
 116732 **msgid** values except the first (including **msgid** values found in an existing dot-po
 116733 file) shall either be commented out or omitted in the resulting dot-po file; if
 116734 omitted, a warning message may be written to standard error. Domain directives
 116735 in the existing dot-po files shall be ignored; the assumption is that all previous

116736 **msgid** values belong to the same domain. The behavior is unspecified if an existing
116737 dot-po file was not created by *xgettext* or has been modified by another application.

116738 **-K** *keyword-spec*

116739 Specify an additional keyword to be looked for:

- 116740 • If *keyword-spec* is an empty string, this shall disable the use of default
116741 keywords for the *gettext* family of functions.
- 116742 • If *keyword-spec* is a C identifier, *xgettext* shall look for strings in the first
116743 argument of each call to the function or macro *keyword-spec*.
- 116744 • If *keyword-spec* is of the form *id:argnum* then *xgettext* shall treat the *argnum*-th
116745 argument of a call to the function or macro *id* as the *msgid* argument, where
116746 *argnum* 1 is the first argument.
- 116747 • If *keyword-spec* is of the form *id:argnum1,argnum2* then *xgettext* shall treat
116748 strings in the *argnum1*-th argument and in the *argnum2*-th argument of a call
116749 to the function or macro *id* as the *msgid* and *msgid_plural* arguments,
116750 respectively.

116751 For all mentioned forms, the application shall ensure that if *argnum2* is given, it is
116752 not equal to *argnum1*. All numeric values shall be converted as specified in item 6
116753 in XBD [Section 12.1](#) (on page 199).

116754 **-n**

116755 Add comment lines to the output file indicating pathnames and line numbers in
116756 the source files where each extracted string is encountered. These lines shall
appear before each **msgid** directive. Such comments should have the format:

116757 `#: pathname1:linenumber1 [pathname2:linenumber2...]`

116758 **-p** *pathname*

116759 Create output files in the directory specified by *pathname* instead of in the current
116760 working directory.

116761 **-x** *exclude-file*

116762 Specify a file containing strings that shall not be extracted from the input files. The
116763 format of *exclude-file* is identical to that of a dot-po file. However, only statements
116764 containing **msgid** directives in *exclude-file* shall be used. All other statements shall
116765 be ignored.

116766 OPERANDS

116767 The following operand shall be supported:

116768 *file* A pathname of an input file containing C-language source code. If '-' is specified
116769 for an instance of *file*, the standard input shall be used.

116770 STDIN

116771 The standard input shall not be used unless a *file* operand is specified as '-'.

116772 INPUT FILES

116773 The input files specified as *file* operands shall be C-language source files. The input file specified
116774 as the option-argument for the **-x** option shall be a dot-po file in the format specified as input for
116775 the *msgfmt* utility.

116776 ENVIRONMENT VARIABLES

116777 The following environment variables shall affect the execution of *xgettext*:

116778	<i>LANG</i>	Provide a default value for the internationalization variables that are unset or null. (See XBD Section 8.2 (on page 157) for the precedence of internationalization variables used to determine the values of locale categories.)
116779		
116780		
116781	XSI <i>LANGUAGE</i>	Determine the location of messages objects if <i>NLSPATH</i> is not set or the evaluation of <i>NLSPATH</i> did not lead to a suitable messages object being found.
116782		
116783	<i>LC_ALL</i>	If set to a non-empty string value, override the values of all the other internationalization variables.
116784		
116785	<i>LC_CTYPE</i>	Determine the locale for the interpretation of sequences of bytes of text data as characters (for example, single-byte as opposed to multi-byte characters in arguments and input files).
116786		
116787		
116788	<i>LC_MESSAGES</i>	
116789		Determine the locale name used to locate messages objects, and the locale that should be used to affect the format and contents of diagnostic messages written to standard error.
116790		
116791		
116792	XSI <i>NLSPATH</i>	Determine the location of messages objects and message catalogs.
116793	ASYNCHRONOUS EVENTS	
116794		Default.
116795	STDOUT	
116796		The standard output shall not be used.
116797	STDERR	
116798		The standard error shall be used for diagnostic messages and may be used for warning messages.
116799		
116800	OUTPUT FILES	
116801		The output files shall be dot-po files in the format specified as input for the <i>msgfmt</i> utility. It is unspecified whether each output file includes a header (msgid "") before the content derived from the input C-language source files.
116802		
116803		
116804	EXTENDED DESCRIPTION	
116805		None.
116806	EXIT STATUS	
116807		The following exit values shall be returned:
116808		0 Successful completion.
116809		>0 An error occurred.
116810	CONSEQUENCES OF ERRORS	
116811		Default.
116812	APPLICATION USAGE	
116813		Implementations differ as to whether they write all output to the default output file or split the output into separate per-domain files. Portable applications can either ensure that each C-language source file contains calls to <i>gettext</i> family functions for only a single domain, or force all output to be to the default output file by using the -K option to override the default keywords.
116814		
116815		
116816		
116817		
116818		Some implementations of <i>xgettext</i> are not able to extract cast strings (unless -a is used), for example casts of literal strings to (const char *). Use of a cast is unnecessary anyway, since the prototypes in <libintl.h> already specify this type.
116819		
116820		

116821 The *xgettext* utility is not required to handle C preprocessor directives. Therefore if, for example,
 116822 calls to *gettext* family functions are wrapped by macros, they might not be found unless the **-K**
 116823 option is used to tell *xgettext* to look for the macro calls.

116824 EXAMPLES

116825 Example 1

116826 The following example shows how **-K** can be used to force all output to be to the default output
 116827 file:

```
116828 xgettext -K "" -K gettext:1 -K dgettext:2 -K dcgettext:2 \  

  116829 -K ngettext:1,2 -K dngettext:2,3 -K dcngettext:2,3 source.c
```

116830 By overriding the default keywords using the **-K** option as above, the *xgettext* utility is directed
 116831 to ignore the *domainname* arguments to the *dgettext()*, *dcgettext()*, *dngettext()*, and *dcngettext()*
 116832 functions. Thus, the utility treats the functions as their respective equivalent without the *d* prefix,
 116833 ignoring the *domainname* argument and writing generated output to the default output file,
 116834 **messages.po**. Additional **-K** options would be needed for the variants of the functions with an
 116835 *_l* suffix if they are used.

116836 Example 2

116837 If the source uses a macro definition such as:

```
116838 #define i18n gettext
```

116839 the use of:

```
116840 xgettext -K i18n:1 source.c
```

116841 will pick up **msgid** values from a line such as:

```
116842 fprintf(stdout, i18n("The value is %s"), value1);
```

116843 RATIONALE

116844 The **-K** option is based on the **-k** option of GNU *xgettext*; the only difference is that GNU's **-k**
 116845 takes an optional option-argument whereas **-K** in this standard has a mandatory option-
 116846 argument in order to comply with the syntax guidelines.

116847 The standard developers considered including functionality equivalent to the **-c**, **-m**, and **-M**
 116848 options in existing implementations. However, those letters could not be used as the syntax
 116849 differed between implementations. The usual solution of adding an uppercase equivalent of
 116850 lowercase options with the standard syntax instead was not possible, for obvious reasons for **-m**
 116851 and **-M**, and as **-C** was already in use for another purpose in one implementation.

116852 The **-s** option is not included as it has been deprecated in at least one implementation because it
 116853 has been found to deprive translators of valuable context.

116854 FUTURE DIRECTIONS

116855 A future version of this standard may change the description of the **-n** option to use "shall"
 116856 instead of "should".

116857 SEE ALSO

116858 [gettext](#), [msgfmt](#)

116859 [XBD Chapter 8](#) (on page 155), [Section 12.2](#) (on page 201)

116860 [XSH gettext](#)

116861 **CHANGE HISTORY**
116862 First released in Issue 8.
116863

129690	POSIX_DEVICE_SPECIFIC_R: Thread-Safe General Terminal	
129691	<i>ttyname_r()</i>	
129692	POSIX_DYNAMIC_LINKING: Dynamic Linking	
129693	<i>dladdr(), dlclose(), dlerror(), dlopen(), dlsym()</i>	
129694	POSIX_FD_MGMT: File Descriptor Management	
129695	<i>dup(), dup2(), dup3(), fcntl(), fgetpos(), fseek(), fseeko(), fsetpos(), ftell(), ftello(), ftruncate(),</i>	
129696	<i>lseek(), rewind()</i>	
129697	POSIX_FIFO: FIFO	
129698	<i>mkfifo()</i>	
129699	POSIX_FIFO_FD: FIFO File Descriptor Routines	
129700	<i>mkfifoat(), mknodat()</i>	
129701	POSIX_FILE_ATTRIBUTES: File Attributes	
129702	<i>chmod(), chown(), fchmod(), fchown(), umask()</i>	
129703	POSIX_FILE_ATTRIBUTES_FD: File Attributes File Descriptor Routines	
129704	<i>fchmodat(), fchownat()</i>	
129705	POSIX_FILE_LOCKING: Thread-Safe Stdio Locking	
129706	<i>flockfile(), frylockfile(), funlockfile(), getc_unlocked(), getchar_unlocked(), putc_unlocked(),</i>	
129707	<i>putchar_unlocked()</i>	
129708	POSIX_FILE_SYSTEM: File System	
129709	<i>access(), chdir(), closedir(), creat(), fchdir(), fpathconf(), fstat(), fstatvfs(), getcwd(), link(),</i>	
129710	<i>mkdir(), mkostemp(), mkstemp(), opendir(), pathconf(), posix_getdents(), readdir(), remove(),</i>	
129711	<i>rename(), rewinddir(), rmdir(), stat(), statvfs(), tmpfile(), tmpnam(), truncate(), unlink()</i>	
129712	POSIX_FILE_SYSTEM_EXT: File System Extensions	
129713	<i>alphasort(), dirfd(), getdelim(), getline(), mkdtemp(), scandir()</i>	
129714	POSIX_FILE_SYSTEM_FD: File System File Descriptor Routines	
129715	<i>faccessat(), fdopendir(), fstatat(), linkat(), mkdirat(), openat(), renameat(), unlinkat(),</i>	
129716	<i>utimensat()</i>	
129717	POSIX_FILE_SYSTEM_GLOB: File System Glob Expansion	
129718	<i>glob(), globfree()</i>	
129719	POSIX_FILE_SYSTEM_R: Thread-Safe File System	
129720	<i>readdir_r()</i>	
129721	POSIX_I18N: Internationalization	
129722	<i>bind_textdomain_codeset(), bindtextdomain(), catclose(), catgets(), catopen(), dcgettext(),</i>	+
129723	<i>dcgettext_l(), dcngettext(), dcngettext_l(), dgettext(), dgettext_l(), dngettext(), dngettext_l(),</i>	+
129724	<i>gettext(), gettext_l(), iconv(), iconv_close(), iconv_open(), ngettext(), ngettext_l(),</i>	
129725	<i>nl_langinfo(), textdomain()</i>	
129726	POSIX_JOB_CONTROL: Job Control	
129727	<i>setpgid(), tcgetpgrp(), tcsetpgrp(), tcgetsid()</i>	
129728	POSIX_MAPPED_FILES: Memory Mapped Files	
129729	<i>mmap(), munmap()</i>	
129730	POSIX_MEMORY_PROTECTION: Memory Protection	
129731	<i>mprotect()</i>	

129820	POSIX_WIDE_CHAR_DEVICE_IO: Device Input and Output	
129821	<i>fgetc()</i> , <i>fgetws()</i> , <i>fputc()</i> , <i>fputws()</i> , <i>fwide()</i> , <i>fwprintf()</i> , <i>fwscanf()</i> , <i>getc()</i> , <i>getwchar()</i> ,	
129822	<i>putc()</i> , <i>putwchar()</i> , <i>ungetc()</i> , <i>vwprintf()</i> , <i>vwscanf()</i> , <i>vwpprintf()</i> , <i>vwsscanf()</i> , <i>wprintf()</i> ,	
129823	<i>wscanf()</i>	
129824	XSI_C_LANG_SUPPORT: XSI General C Library	
129825	<i>a64l()</i> , <i>daylight</i> , <i>drand48()</i> , <i>erand48()</i> , <i>ffs()</i> , <i>ffsl()</i> , <i>ffsll()</i> , <i>getdate()</i> , <i>hcreate()</i> , <i>hdestroy()</i> ,	
129826	<i>hsearch()</i> , <i>initstate()</i> , <i>insque()</i> , <i>rand48()</i> , <i>l64a()</i> , <i>lcong48()</i> , <i>lfind()</i> , <i>lrand48()</i> , <i>lsearch()</i> ,	
129827	<i>memcpy()</i> , <i>mrnd48()</i> , <i>nrnd48()</i> , <i>random()</i> , <i>remque()</i> , <i>seed48()</i> , <i>setstate()</i> , <i>siggam</i> ,	
129828	<i>srand48()</i> , <i>srandom()</i> , <i>strptime()</i> , <i>swab()</i> , <i>tdelete()</i> , <i>tfind()</i> , <i>timezone</i> , <i>tsearch()</i> , <i>twalk()</i>	
129829	XSI_DBM: XSI Database Management	
129830	<i>dbm_clearerr()</i> , <i>dbm_close()</i> , <i>dbm_delete()</i> , <i>dbm_error()</i> , <i>dbm_fetch()</i> , <i>dbm_firstkey()</i> ,	
129831	<i>dbm_nextkey()</i> , <i>dbm_open()</i> , <i>dbm_store()</i>	
129832	XSI_DEVICE_IO: XSI Device Input and Output	
129833	<i>fntmsg()</i> , <i>readv()</i> , <i>writv()</i>	
129834	XSI_DEVICE_SPECIFIC: XSI General Terminal	
129835	<i>grantpt()</i> , <i>posix_openpt()</i> , <i>ptsname()</i> , <i>unlockpt()</i>	
129836	XSI_FILE_SYSTEM: XSI File System	
129837	<i>basename()</i> , <i>dirname()</i> , <i>lockf()</i> , <i>mknod()</i> , <i>nftw()</i> , <i>realpath()</i> , <i>seekdir()</i> , <i>sync()</i> , <i>telldir()</i>	
129838	XSI_GENERAL_TERMINAL_R: XSI Thread-Safe General Terminal	
129839	<i>ptsname_r()</i>	
129840	XSI_IPC: XSI Interprocess Communication	
129841	<i>ftok()</i> , <i>msgctl()</i> , <i>msgget()</i> , <i>msgrcv()</i> , <i>msgsnd()</i> , <i>semctl()</i> , <i>semget()</i> , <i>semop()</i> , <i>shmat()</i> , <i>shmctl()</i> ,	
129842	<i>shmdt()</i> , <i>shmget()</i>	
129843	XSI_MATH: XSI Maths Library	
129844	<i>j0()</i> , <i>j1()</i> , <i>jn()</i> , <i>y0()</i> , <i>y1()</i> , <i>yn()</i>	
129845	XSI_MULTI_PROCESS: XSI Multiple Process	
129846	<i>getpriority()</i> , <i>getrlimit()</i> , <i>getrusage()</i> , <i>nice()</i> , <i>setpriority()</i> , <i>setrlimit()</i>	
129847	XSI_SIGNALS: XSI Signal	
129848	<i>killpg()</i> , <i>sigaltstack()</i>	
129849	XSI_SINGLE_PROCESS: XSI Single Process	
129850	<i>gethostid()</i> , <i>putenv()</i>	
129851	XSI_SYSTEM_DATABASE: XSI System Database	
129852	<i>endpwent()</i> , <i>getpwent()</i> , <i>setpwent()</i>	
129853	XSI_SYSTEM_LOGGING: XSI System Logging	
129854	<i>closelog()</i> , <i>openlog()</i> , <i>setlogmask()</i> , <i>syslog()</i>	
129855	XSI_THREADS_EXT: XSI Threads Extensions	
129856	<i>pthread_attr_getstack()</i> , <i>pthread_attr_setstack()</i>	
129857	XSI_USER_GROUPS: XSI User and Group	
129858	<i>endgrent()</i> , <i>endutxent()</i> , <i>getgrent()</i> , <i>getresgid()</i> , <i>getresuid()</i> , <i>getutxent()</i> , <i>getutxid()</i> ,	+
129859	<i>getutxline()</i> , <i>pututxline()</i> , <i>setgrent()</i> , <i>setregid()</i> , <i>setresgid()</i> , <i>setresuid()</i> , <i>setreuid()</i> , <i>setutxent()</i>	+
129860	XSI_WIDE_CHAR: XSI Wide-Character Library	
129861	<i>wcswidth()</i> , <i>wcwidth()</i>	